The iMagine AI platform

WP4 update and status

Álvaro López García - Advanced Computing and e-Science Group (IFCA-CSIC-UC)
(aloga@ifca.unican.es)

iMagine RP1 review

December 5th 2023

DEEP-HDC, AI4EOSC, iMagine, AI4OS...

- **DEEP-1**, **DEEP-2**: Platform releases
- Platform and software tightly coupled and interlinked, difficult to self-deploy and customize

- **AI4EOSC platform** → Platform “powered by AIOS”
  - DEEP-3 → AI4EOSC-3
- **AI4OS** → software distribution
  - Possible to build custom platforms, partially integrated with AI4EOSC platform (i.e. reusing services) or not
  - https://github.com/AI4OS

- **iMagine AI platform**
  - Customized platform for AI image processing
  - Support for AI service deployment and creation
  - Exploitation if DEEP/AI4OS software as technology provider

DEEP-HDC, AI4EOSC, iMagine, AI4OS…

- **DEEP-1**, **DEEP-2**: Platform releases
- Platform and software tightly coupled and interlinked, difficult to self-deploy and customize

(large) User input, missing features

- **AI4EOSC platform** → Platform “powered by AIOS”
  - DEEP-3 → AI4EOSC-3
- **AI4OS** → software distribution
  - Possible to build custom platforms, partially integrated with AI4EOSC platform (i.e. reusing services) or not
  - https://github.com/AI4OS

Feature delivery

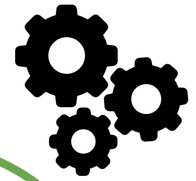

- **iMagine AI platform**
  - Customized platform for AI image processing
  - Support for AI service deployment and creation
  - Exploitation if DEEP/AI4OS software as technology provider

iMagine AI platform customizations

A distributed and federated platform

- Integration of different cloud computing providers from the EGI Federated Cloud
  - GPU access
    - CSIC Scientific Cloud (es), INCD Cloud (pt), TÜBITAK-ULAKBIM (tk)
  - CPU access
    - Walton Institute (ie)

- Transparent access to pan-EU e-Infrastructures state of the art resources
  - Distributed access transparent for users
  - Integration with EOSC ecosystem

- Not imposing any special libraries or hard requirements for users
  - Only requirement is (light) integration with API

- Compatible with KServe community API for model delivery (in progress)

User engagement: how we deliver?

- EOSC portal onboarding
  - https://marketplace.eosc-portal.eu/services/imaging-ai-platform-for-aquatic-science
- Webinars, user meetings
  - iMagine Competence Centre workshop: https://indico.egi.eu/event/5999/
  - EGI Conference 2023: https://whova.com/web/M8zkrnLo5DUwlnug54VINPkHTdssyI49PHa20jCW2Qg%3D/Agenda/
  - AI4EOSC + iMagine user workshop: https://indico.scc.kit.edu/event/3845/
- Comprehensive documentation
  - https://docs.ai4os.eu/
- Customized user guide and quickstart
  - https://confluence.egi.eu/display/IMPAIP/User+guide

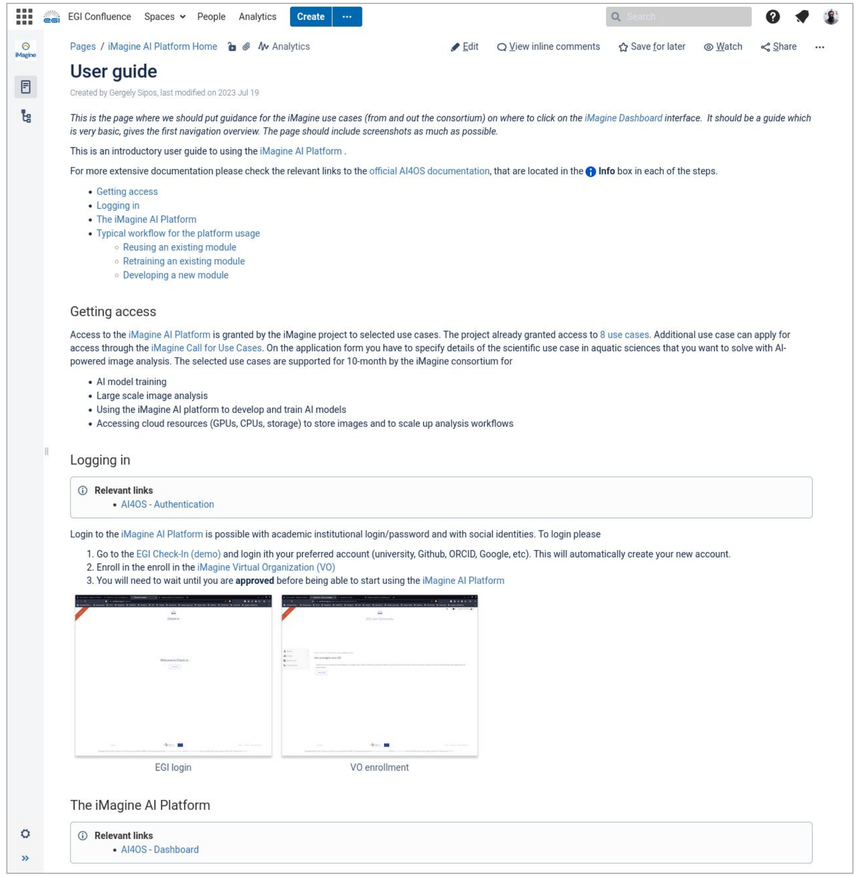

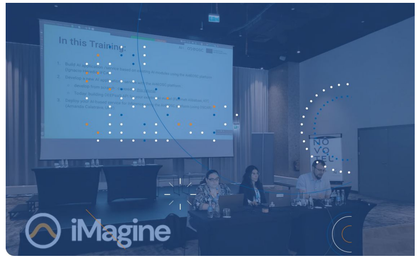

iMagine

Platform status

Integration status

Federated Compute Infrastructure (T4.3)

- All sites now supporting iMagine VO (Operators group) to deploy platform services
- Usage reported through EGI accounting system (link)

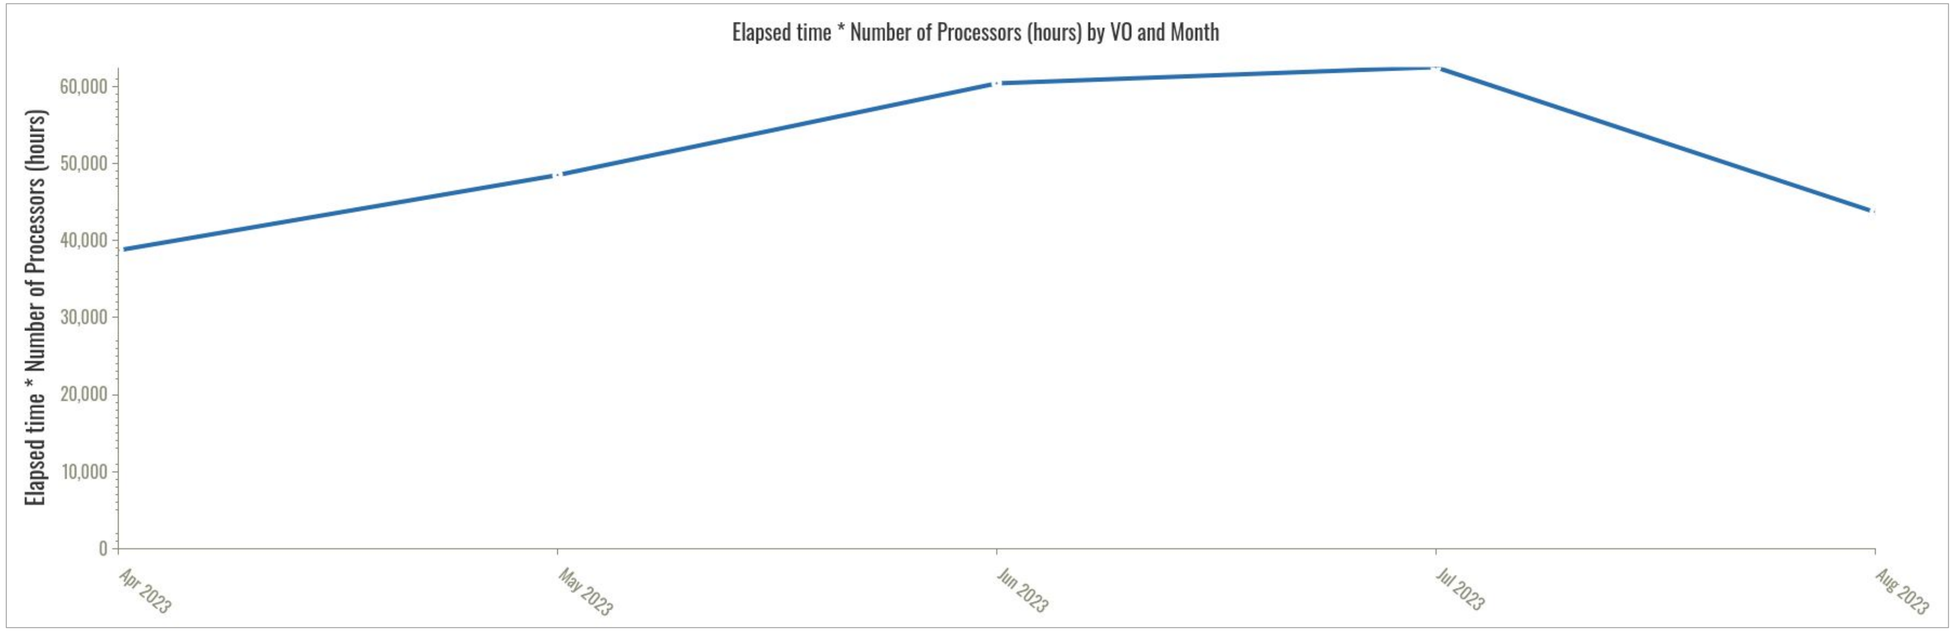

Integration status

Federated Compute Infrastructure (T4.3)

- All sites now supporting iMagine VO (Operators group) to deploy platform services
- Usage reported through EGI accounting system (link)

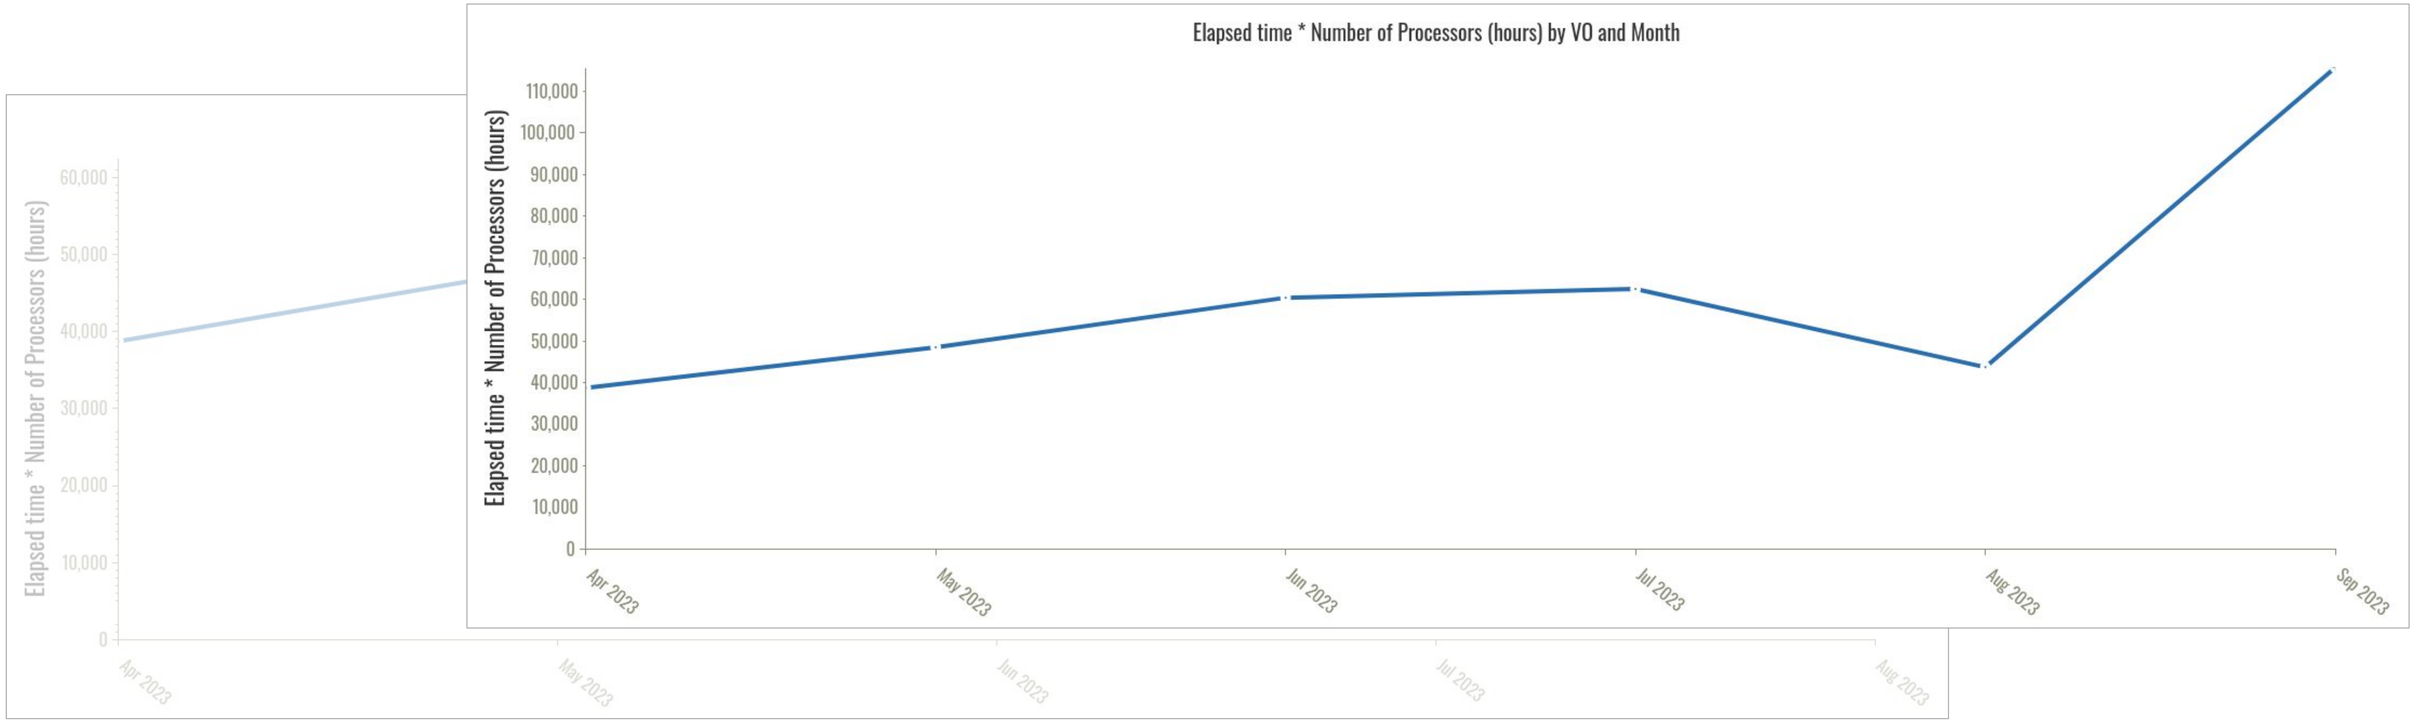

Integration status

*Imagine AI **development** platform (T4.1)*

- Production system transitioned from DEEP software stack to AI4OS
  - Now running fully in AI4OS
  - EGI CheckIn integration
- Production system
  - CSIC, INCD and TUBITAK (GPU)
  - Integrating Walton (CPU)
- Development testbed
  - CSIC, INCD

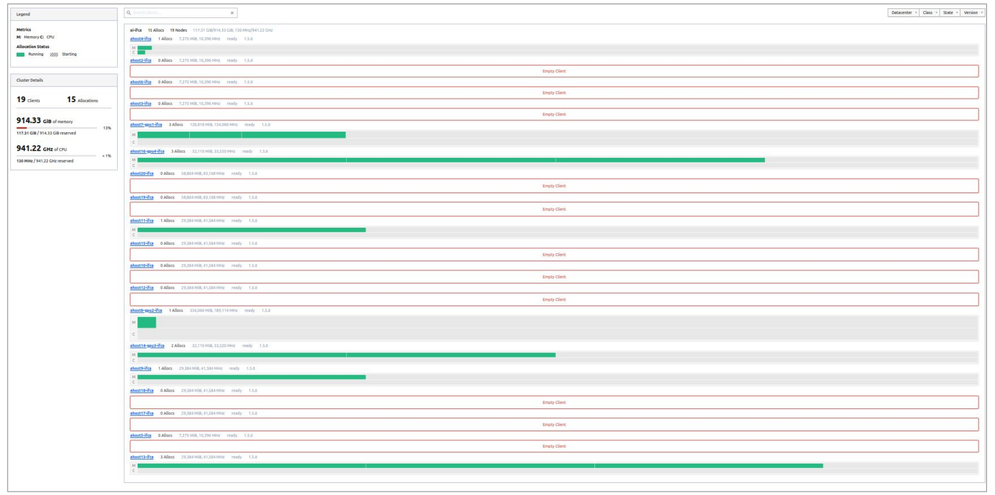

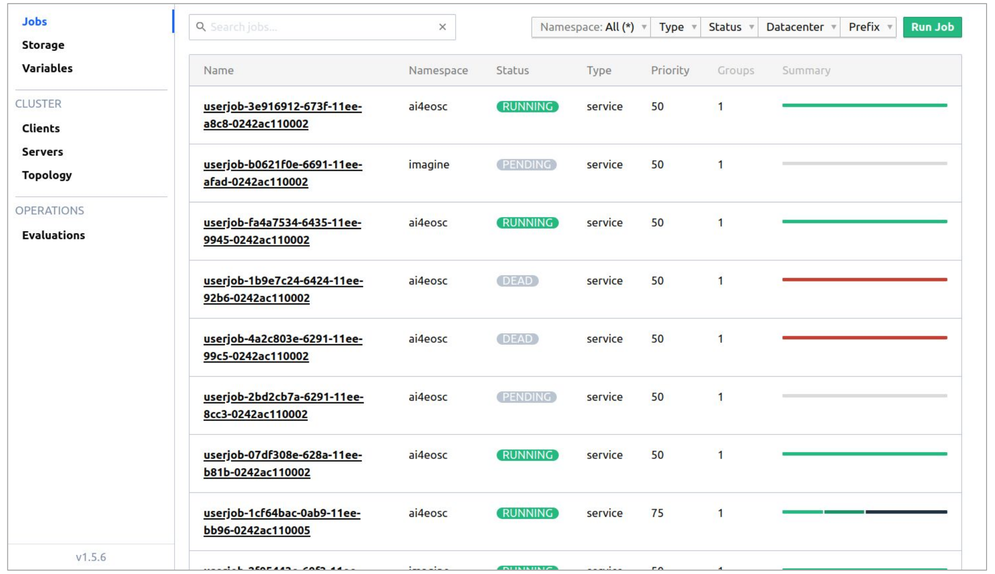

Integration status

*Imagine AI **service** platform (T4.2)*

- Preliminary integration work, bulk work not yet started
  - Depending on WP5 work to deliver
- Exploiting CPU resources for inference
- Different approaches for deployment of services
  - Standalone (i.e. single server in a VM through EGI IM)
  - OSCAR clusters (Kubernetes based, serverless inference, through EGI IM)
  - Platform-level preview services (self-contained in platform)

iMagine

For end users and user communities

Platform features

iMagine

AI/ML application development lifecycle

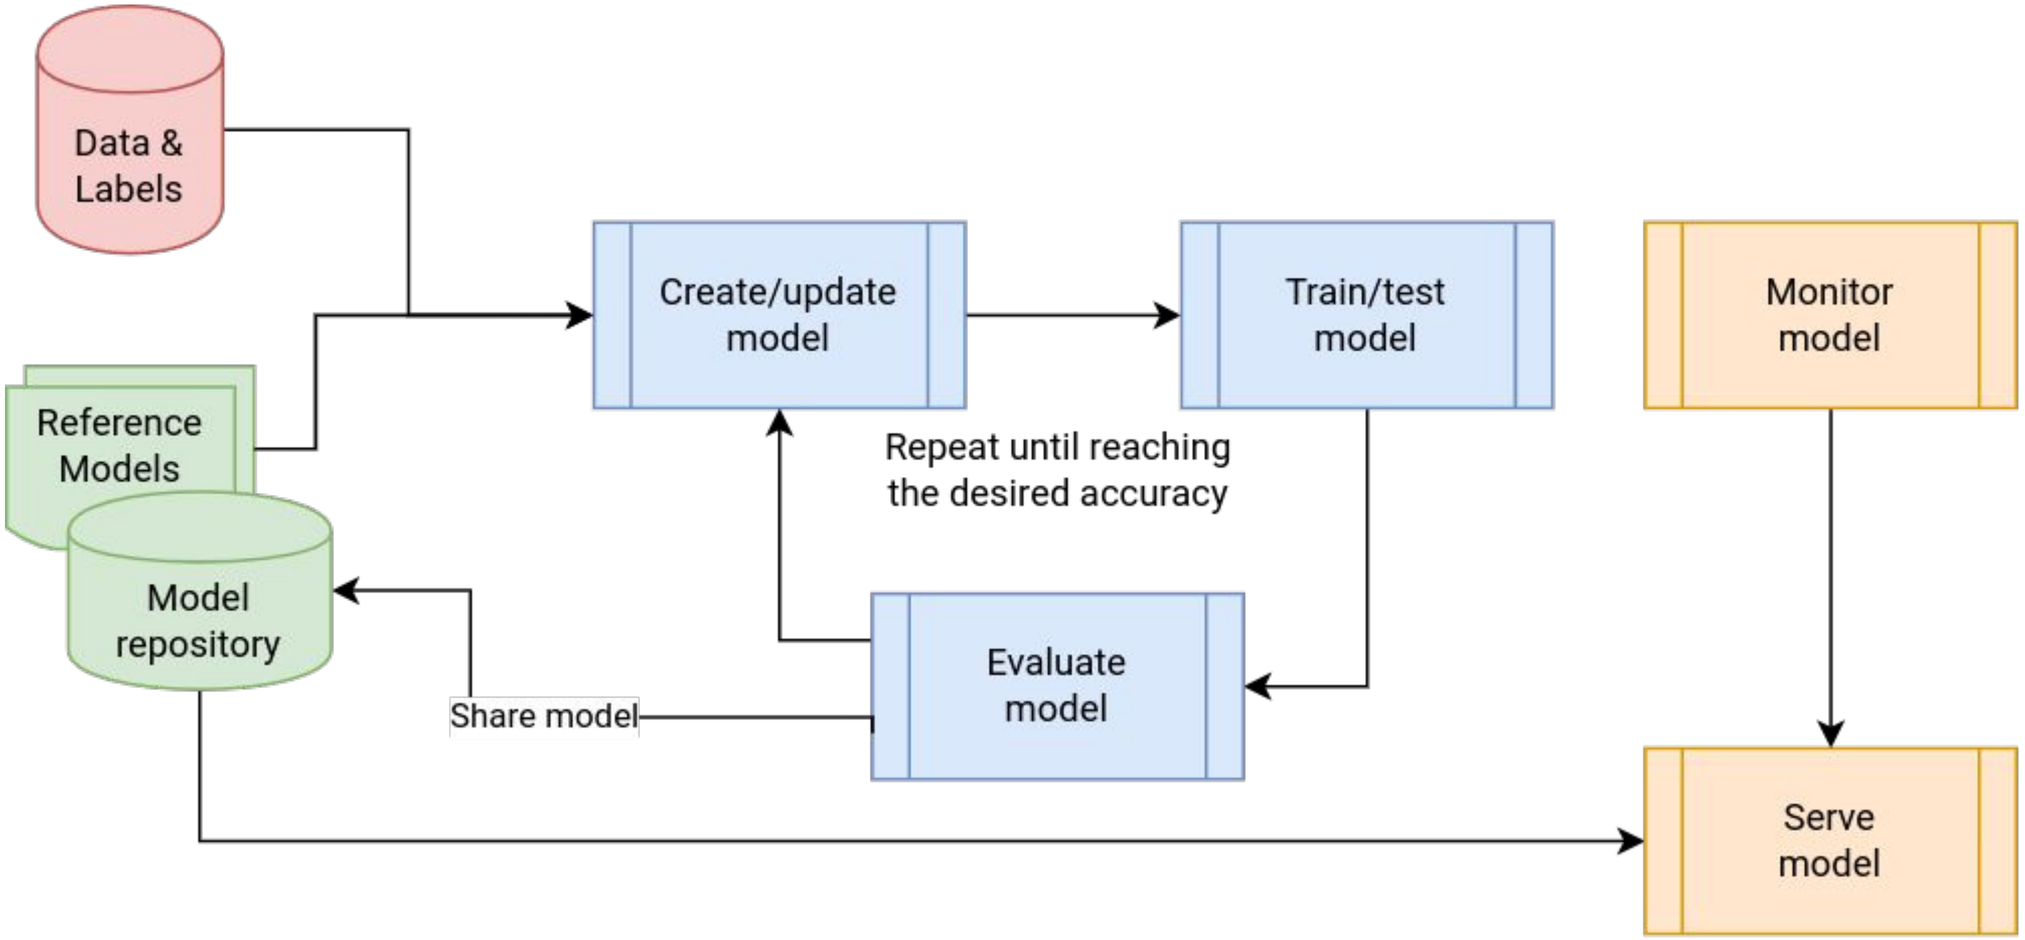

Services for AI/ML development

Model marketplace

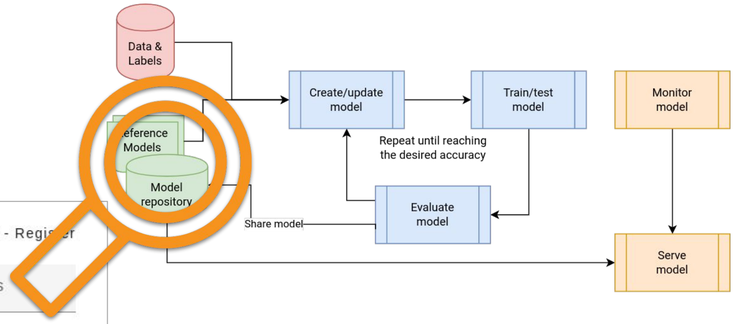

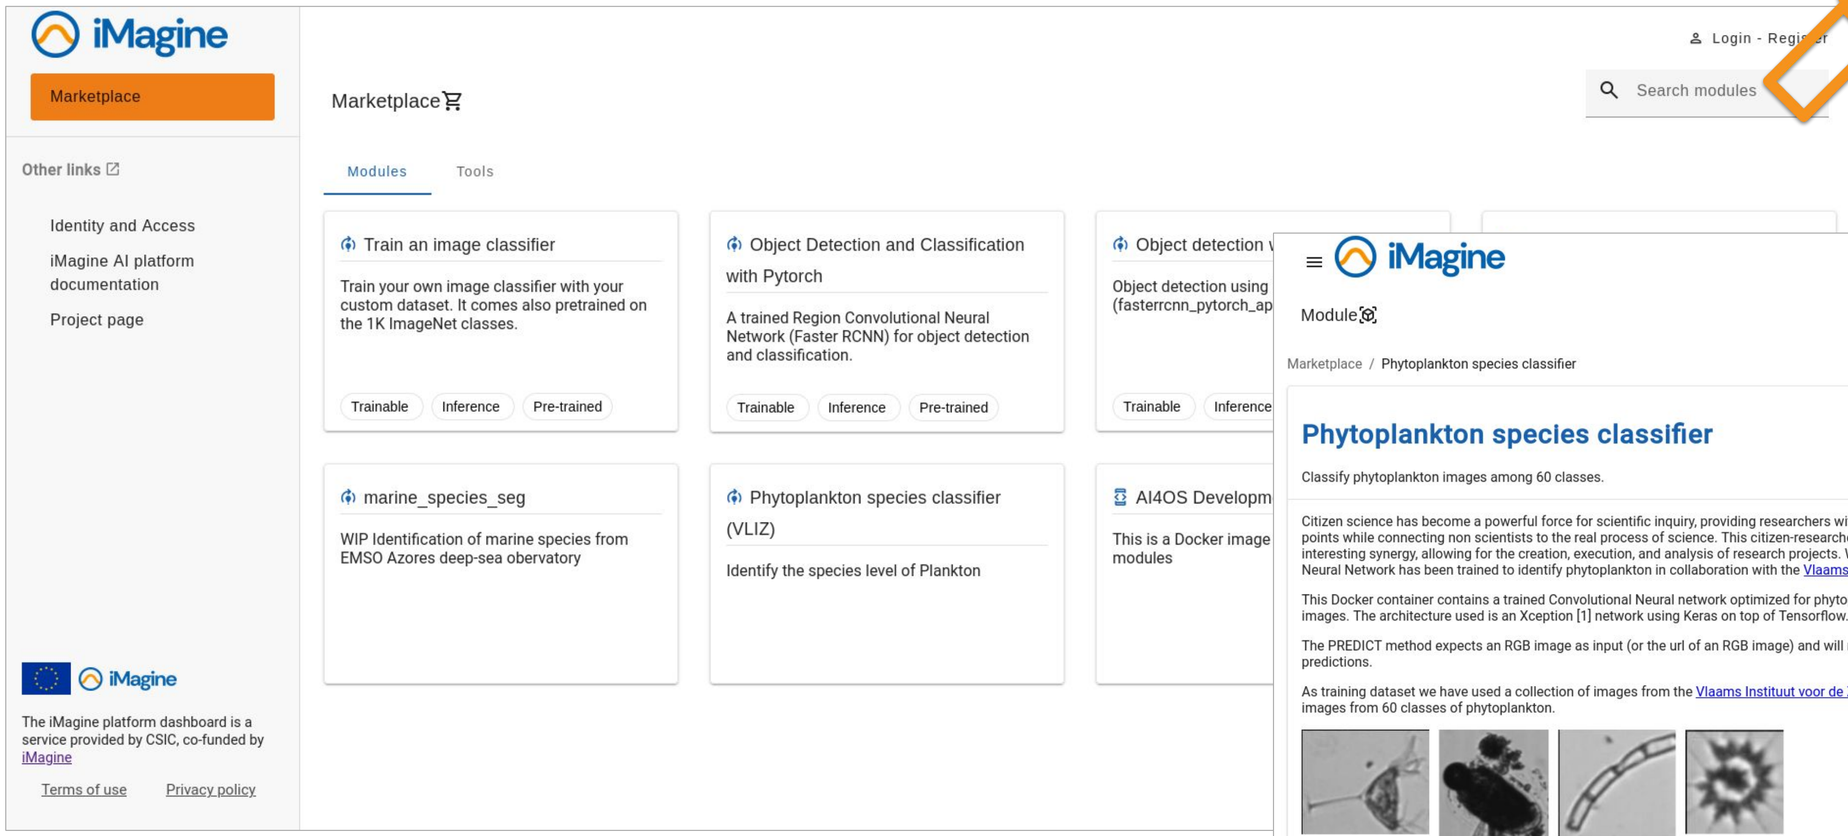

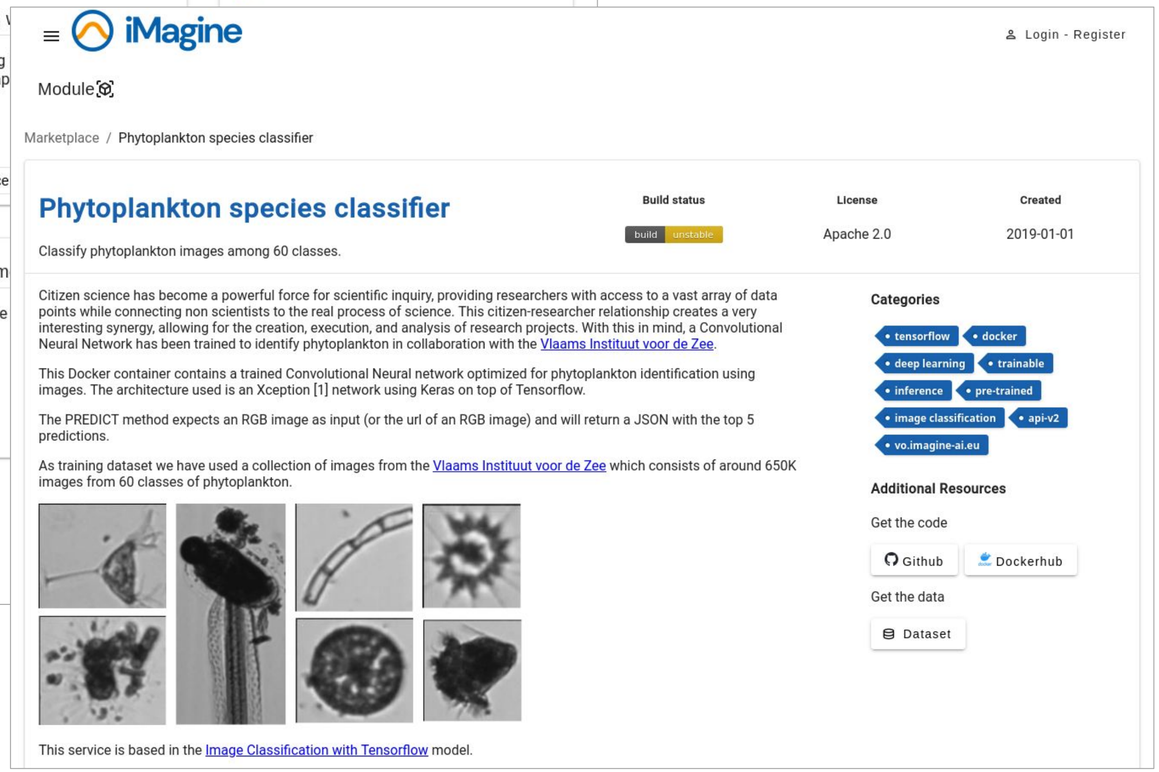

https://dashboard.cloud.imagine-ai.eu/marketplace

https://marketplace.eosc-portal.eu/services/imaging-ai-platform-for-aquatic-science

Services for AI/ML development

Dev tools: CVAT annotation

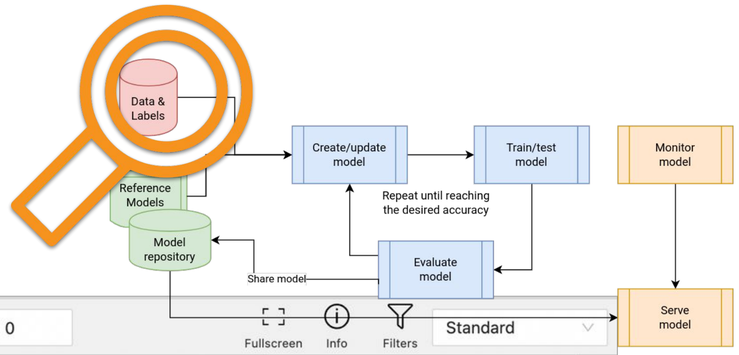

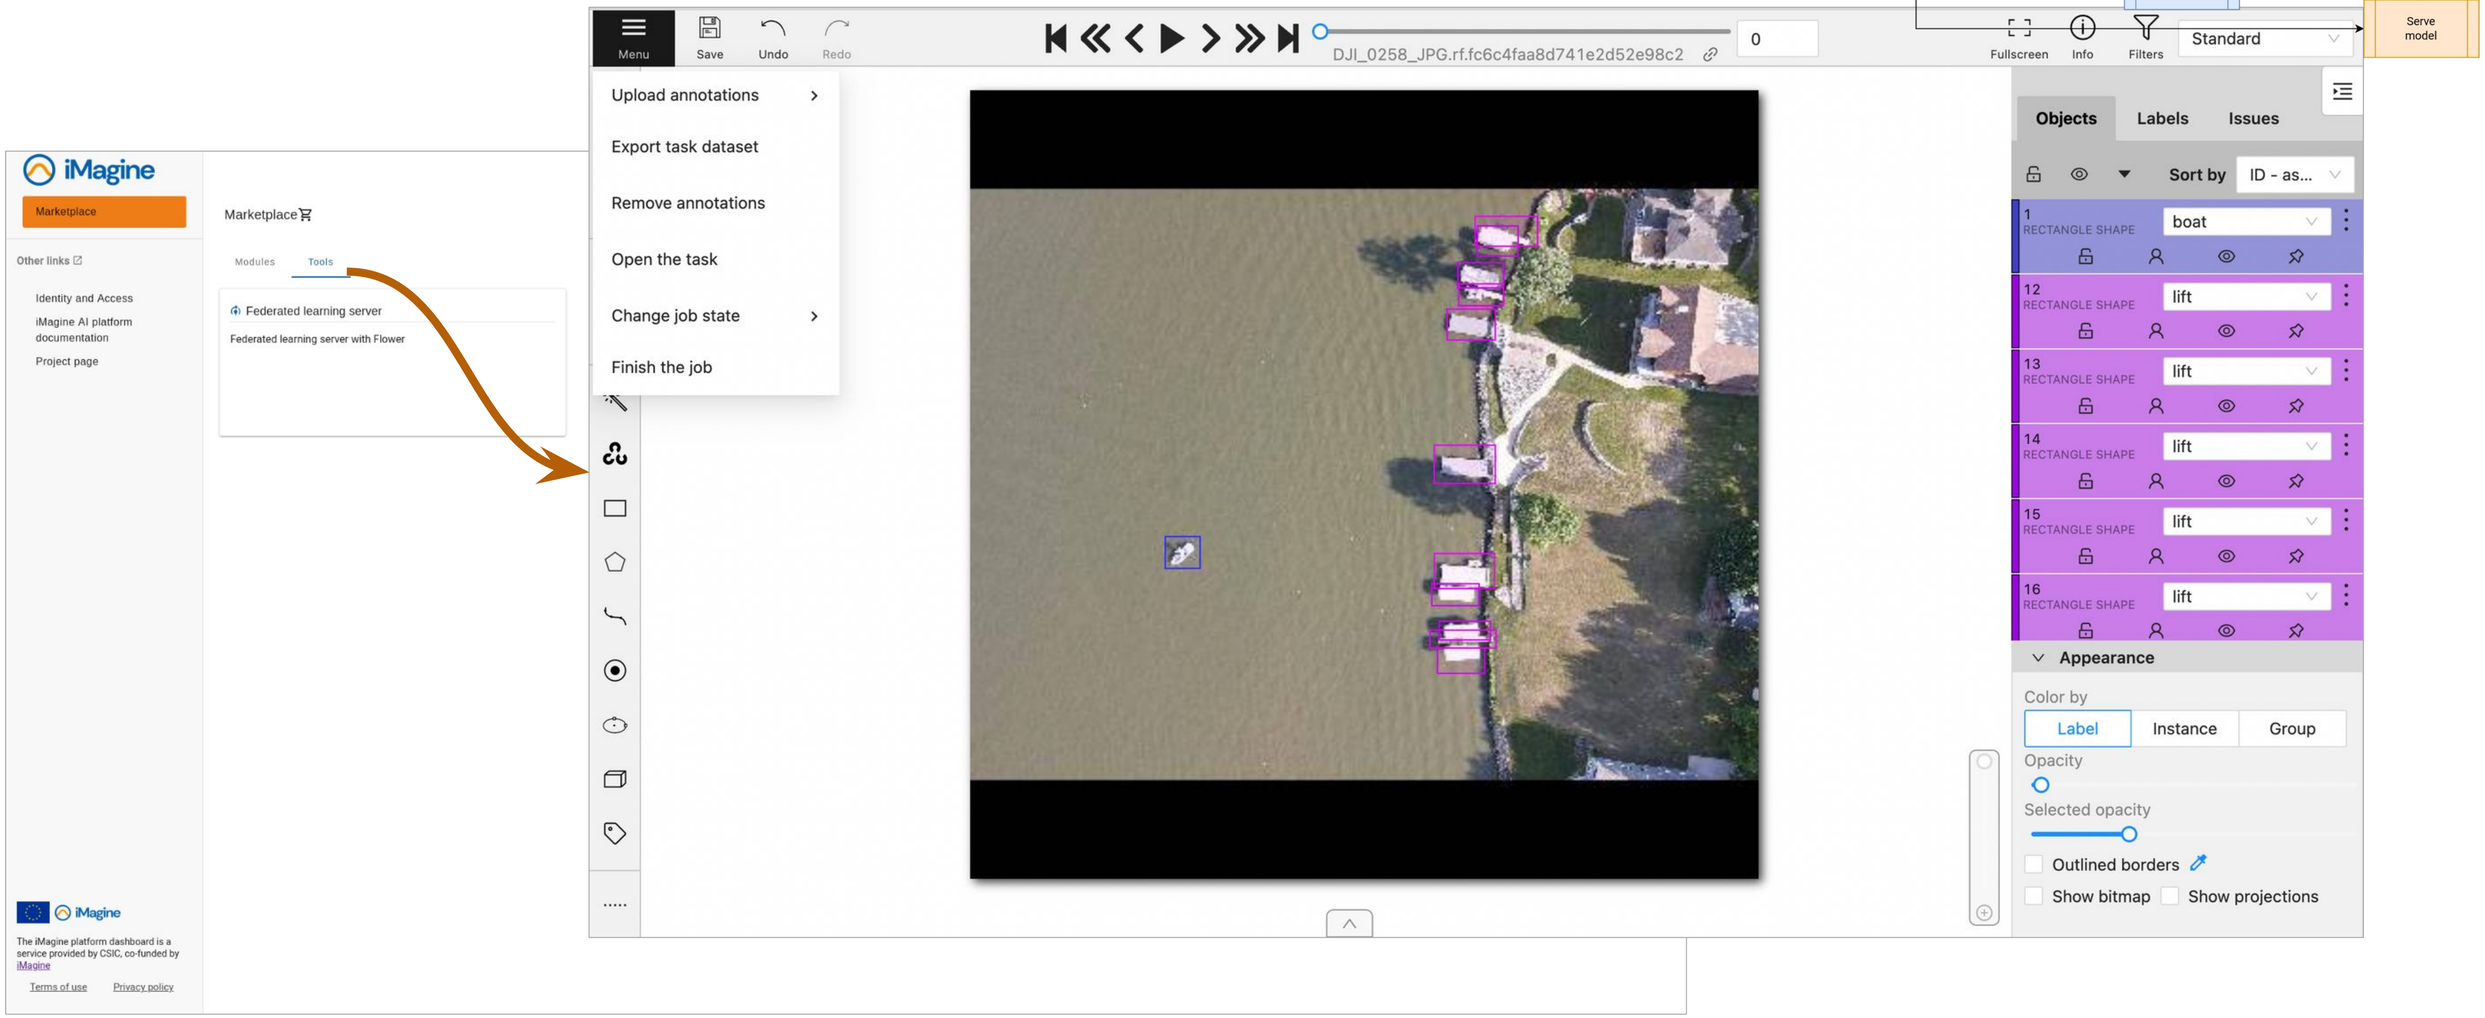

Services for AI/ML development

Dev tools: sandbox and online IDE

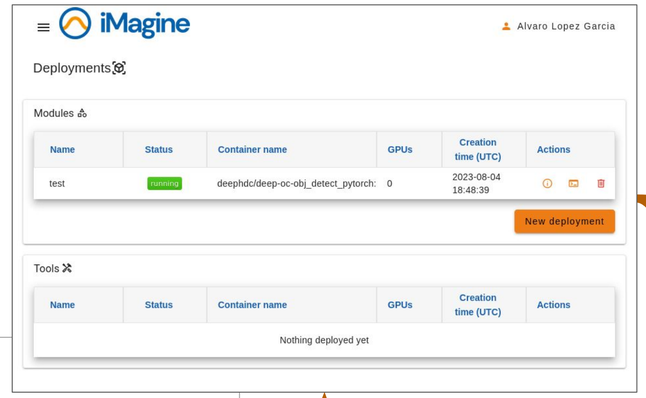

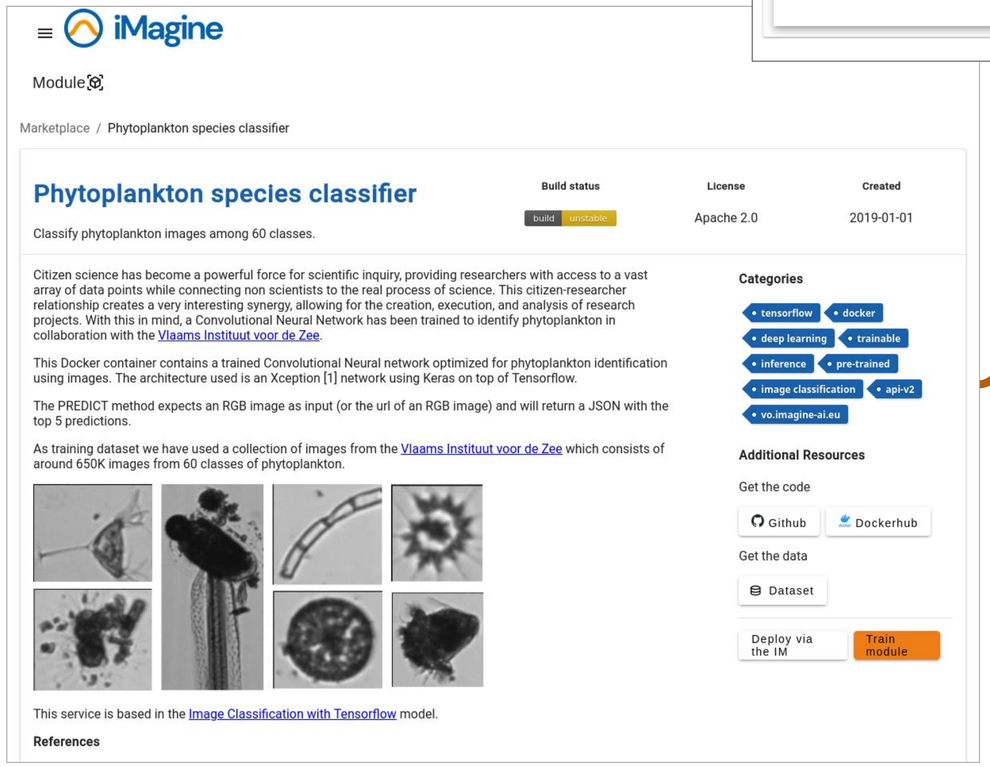

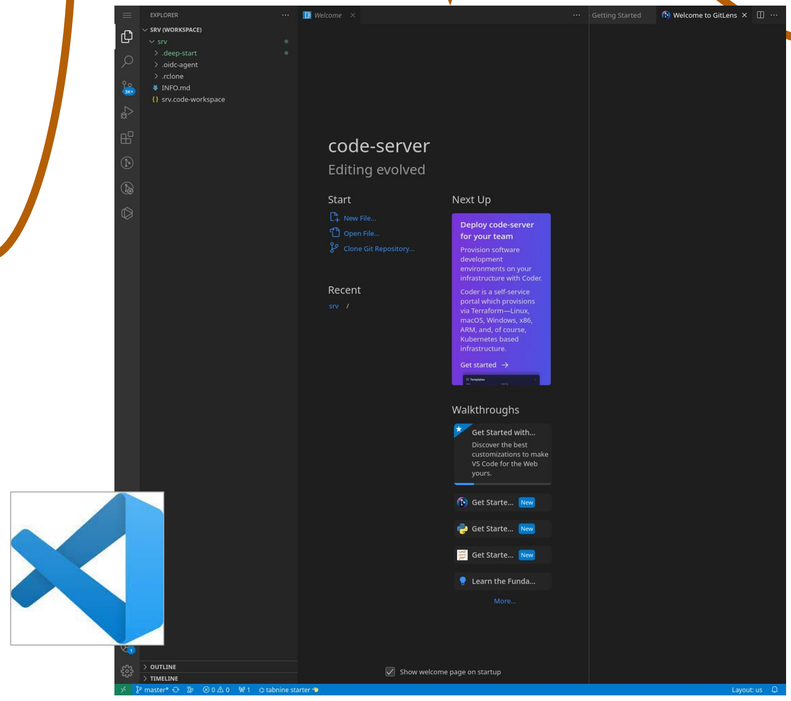

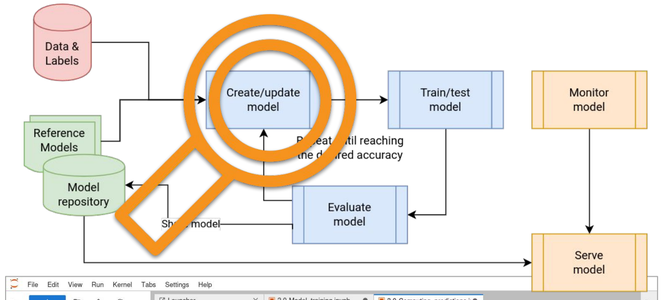

Services for AI/ML development

Training: Transparent GPU access

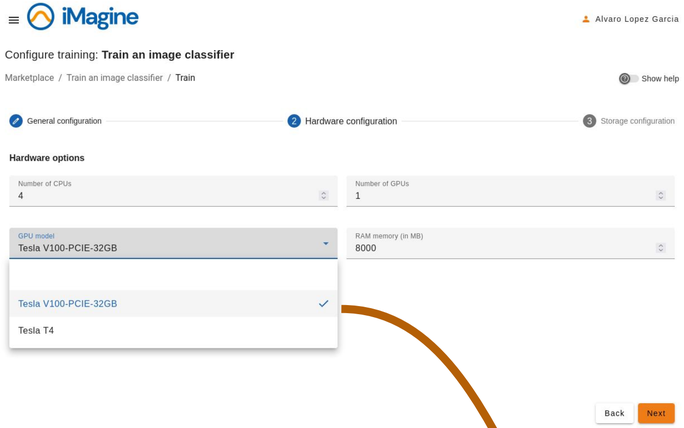

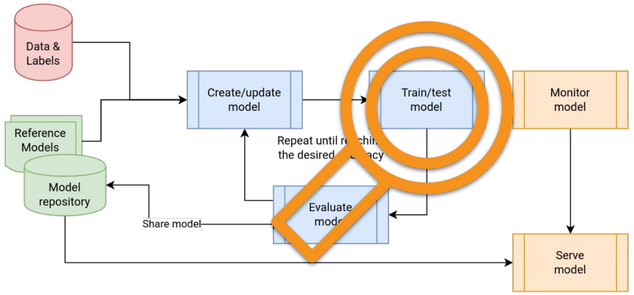

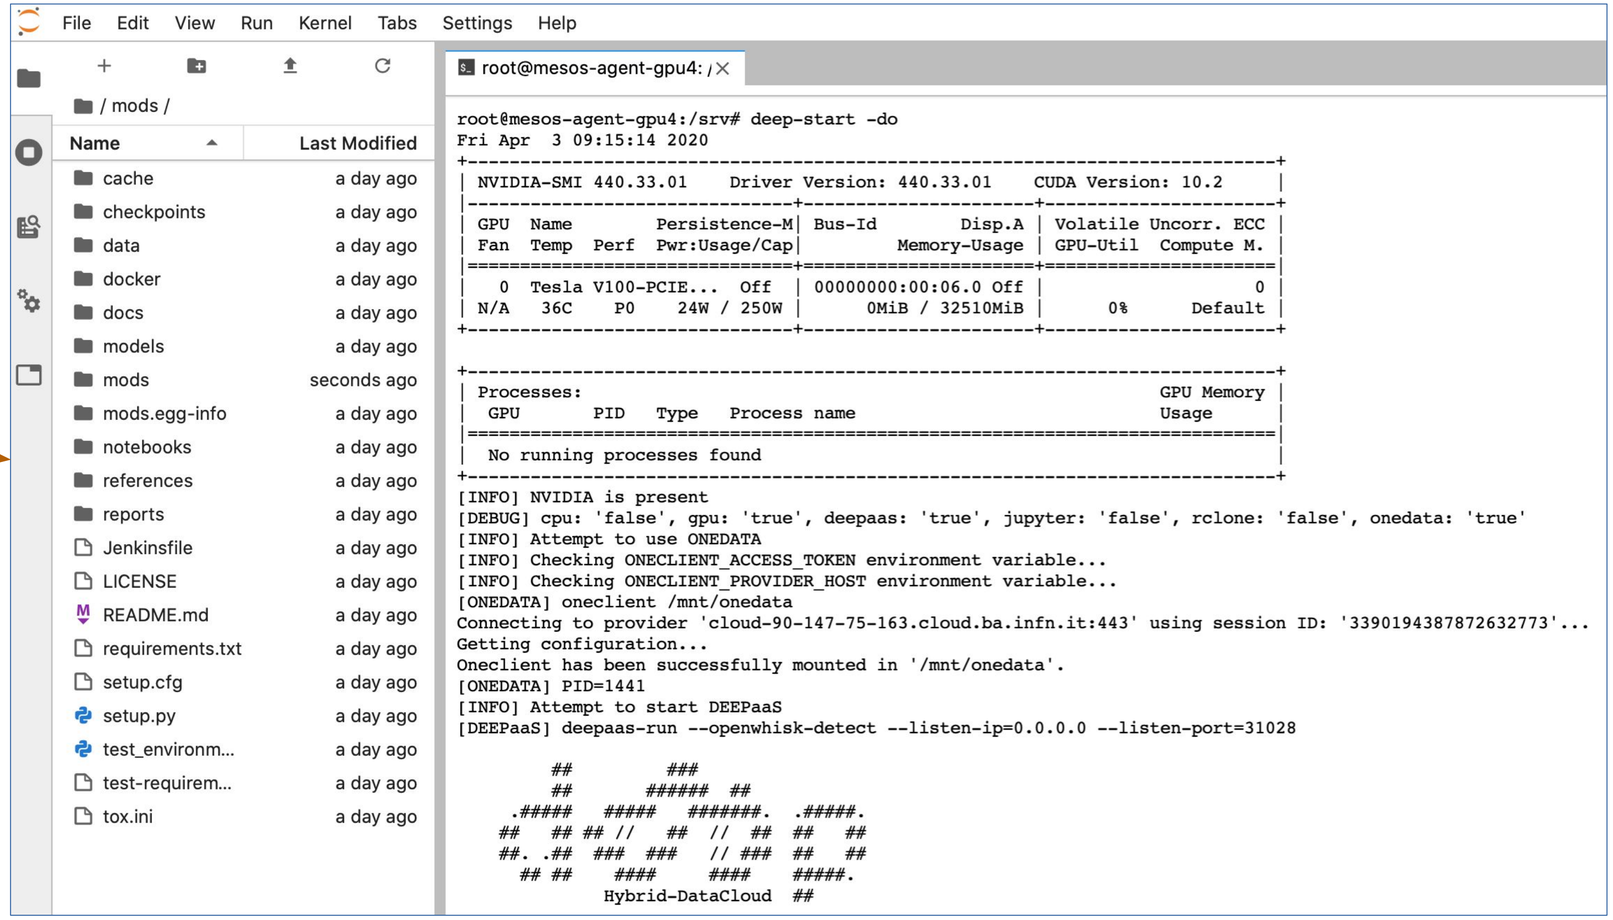

Services for AI/ML development

Training: Transparent GPU access

- Automatic platform tasks to collect information from running jobs, then aggregation is performed
- Time series delivered directly by API
- Integration in upstream dashboard is in progress

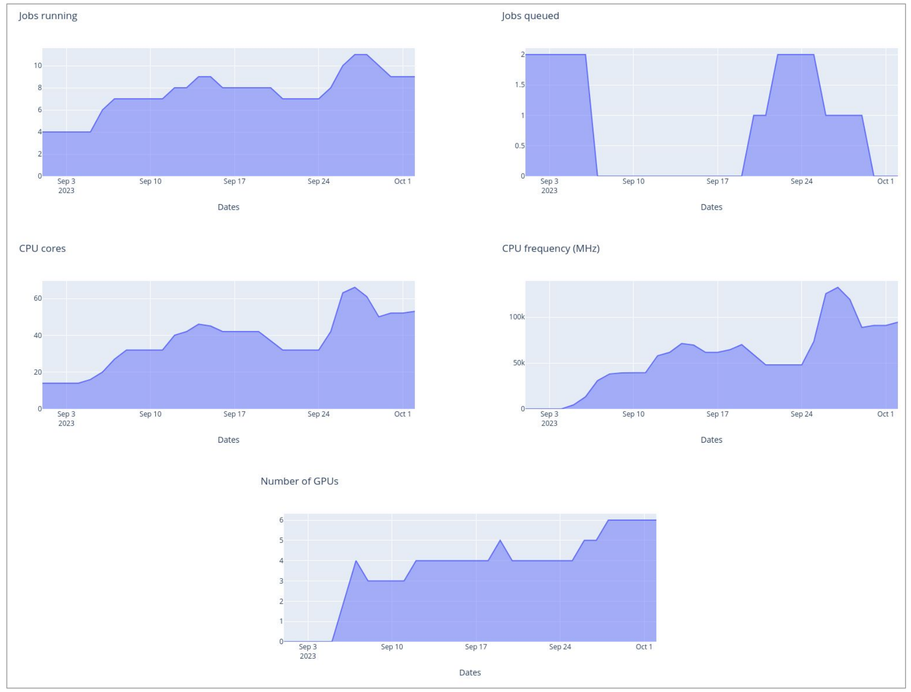

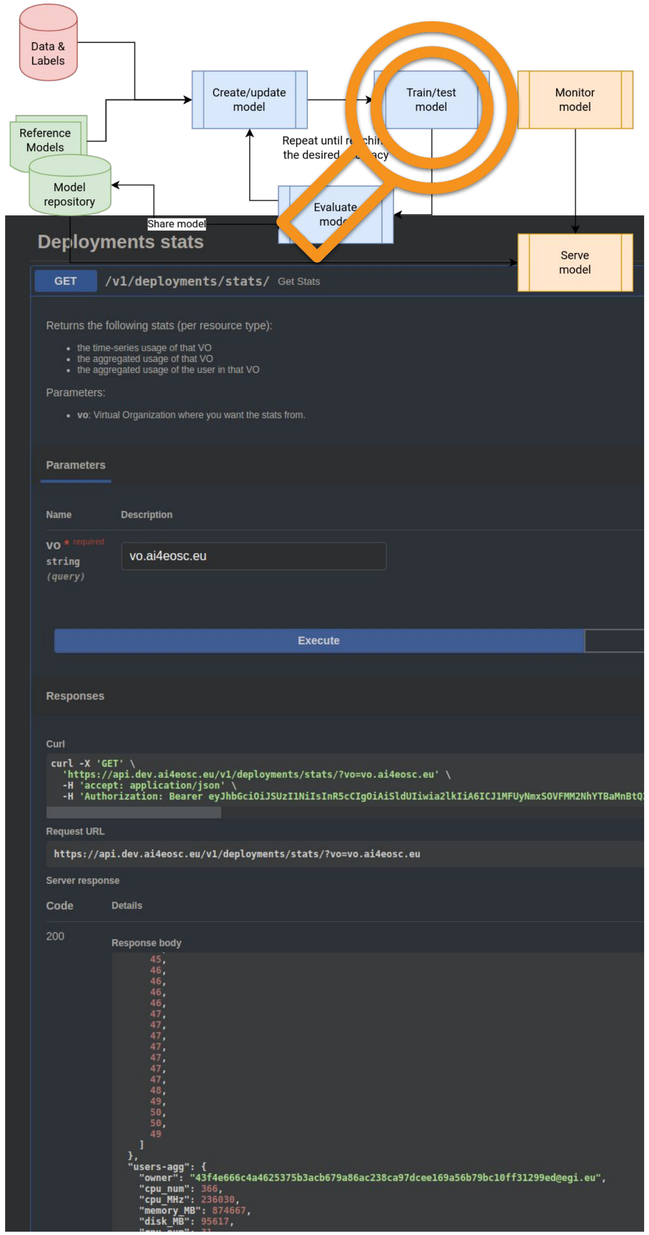

Services for AI/ML development

Deploy and monitor

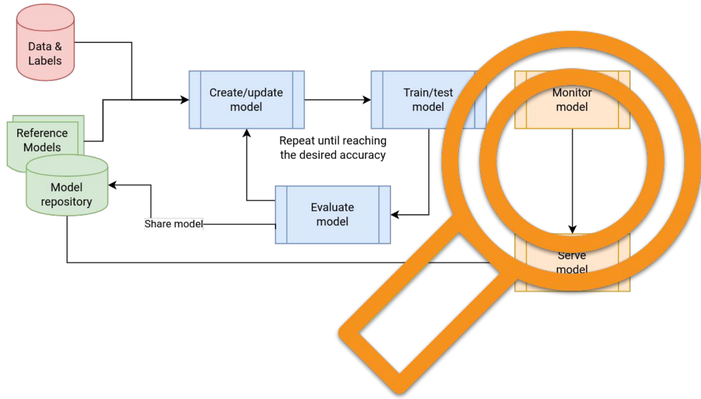

- Deployment
  - Providing services for deployment as services:
    - Through IM an a different cloud (done)
    - Through OSCAR in a platform-managed cluster (partially done)
    - Through OSCAR in user own resources (in progress)
    - Through AI4-PAPI in iMagine AI platform resources (planned)
  - Composite-AI tools (visually build more complex models) (in progress)
- MLOps
  - CI/CD for ML development, deployment, monitoring and operations
- Monitoring
  - Tools to instrument ML models in production (i.e. drift detection)

Services for AI/ML development

Deployment: standalone service

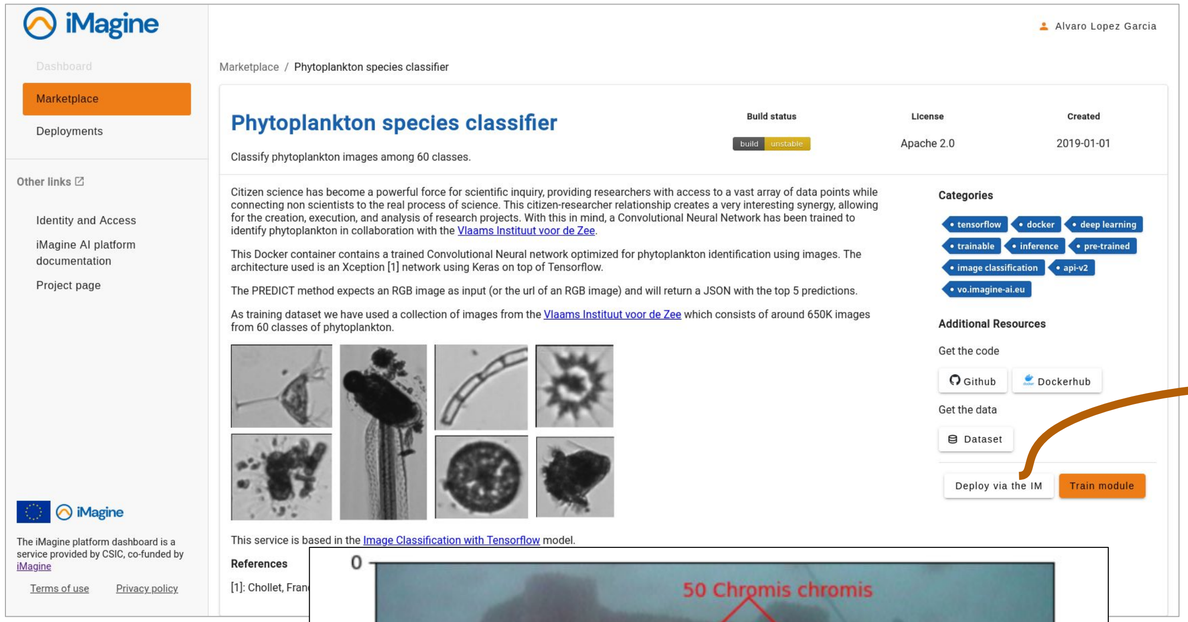

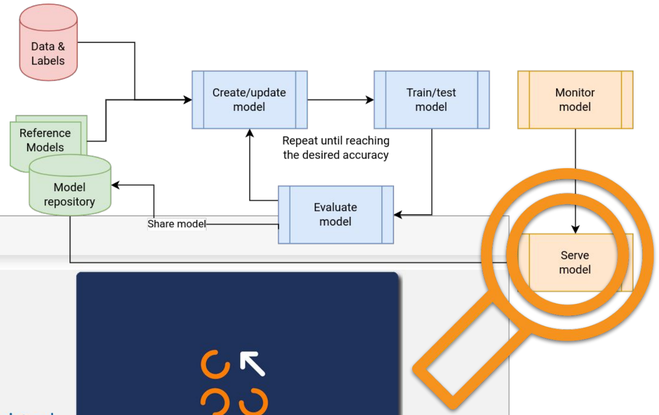

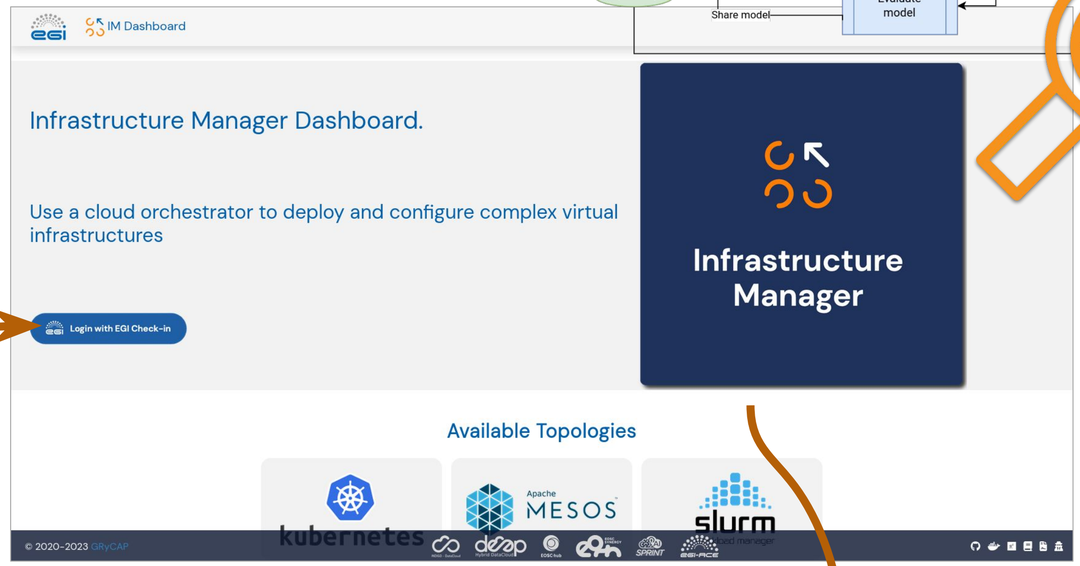

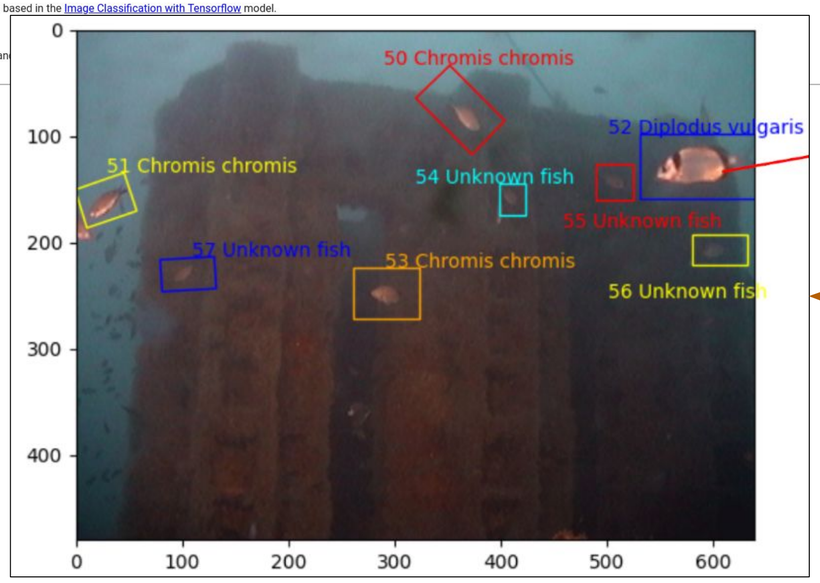

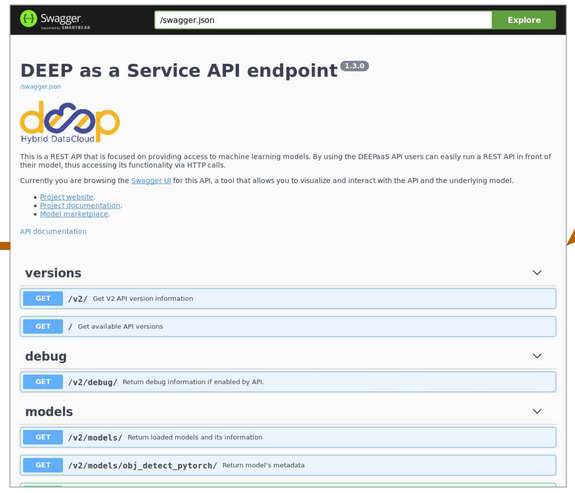

Services for AI/ML development

Deployment: Composite AI

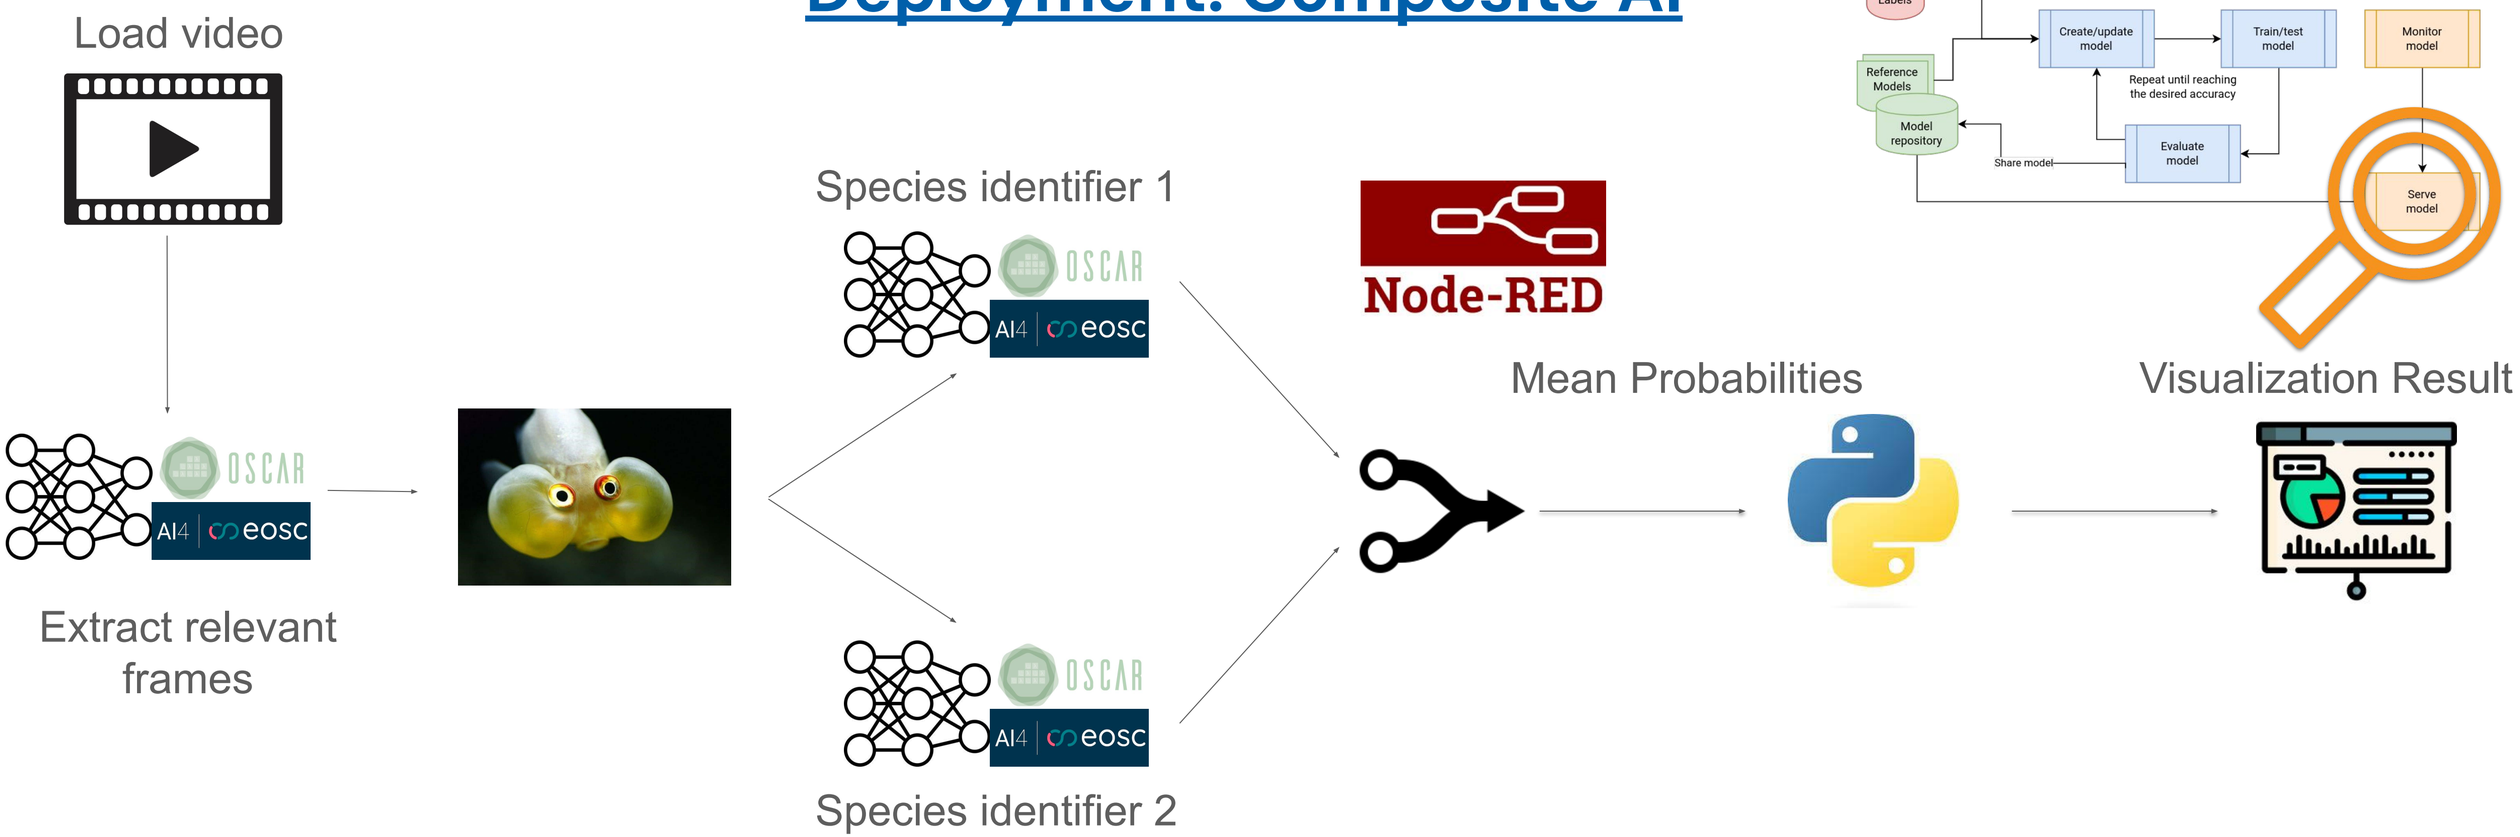

- Use case: multiple AI models can be triggered for inference and later aggregate the results for enhanced accuracy
- Reuse functions (subflow)
- Visual support (drag & drop + customization)
- Minimize orchestration costs

Services for AI/ML development

Deployment: Drift detection

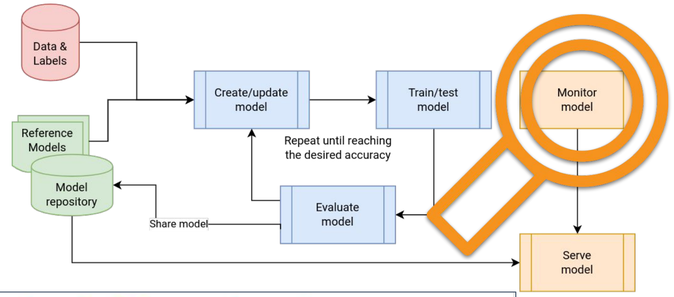

- Monitoring of models in production is not enough
  - Model learns from data, data is not stationary
  - Concept learnt by them model may change over time
- Data and concept drift detection → essential to build more robust models
- Frouros: state-of-the-art library for drift detection in ML problems
  - https://github.com/IFCA/frouros
- Ongoing work towards online services for drift detection → MLOps pipelines with drift detection

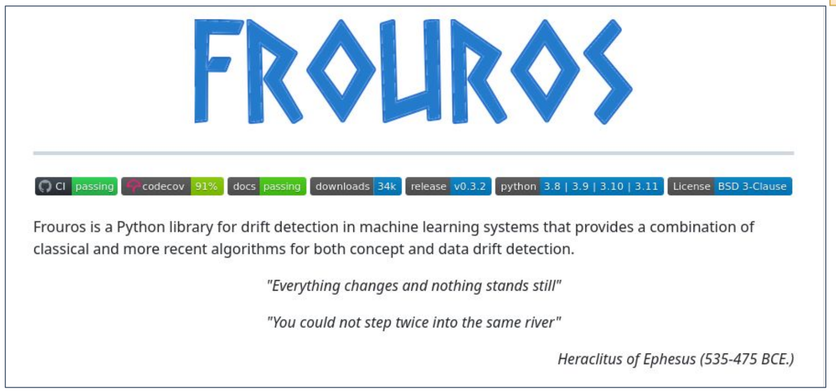

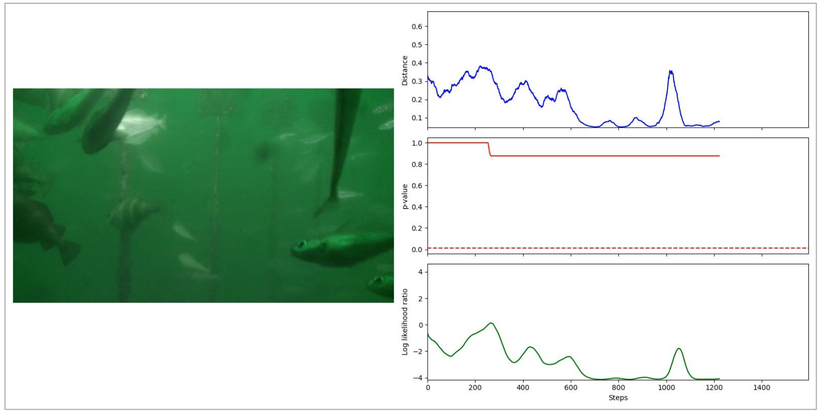

Example: data drift detection in underwater video

iMagine

Conclusions, next steps

Wrap-up

iMagine

Next features

Expected (user) features for 2024

- Final delivery of resource consumption accounting in the dashboard
  - Improving GPU usage and release
  - Work in progress for preemptible, interactive jobs
- Streamlined model deployment as services
  - Including initial MLOps pipelines
- Distributed learning schemes
  - Horovod and Tensorflow parameter server
- Experiment centric dashboard (Q2 2024)
  - Integration of related tools (CVAT), tracking (MLFlow), training deployments and service deployments in a single space
- Integrating metadata schemes in the model registry and dashboard
  - Both generic for ML and domain-specific
  - FAIRness of models and ML assets

Thank you!

imagine-ai.eu

IiMagine receives funding from the European Union's Horizon Europe research and innovation programme under grant agreement No. 101058625

iMagine

Backup slides

Background, ecosystem, collaborations

2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	2026	…

DEEP 3 → AI4EOSC 3 platform
AI4OS → platform agnostic software stack

DEEP Platform and services

INDIGO-DataCloud
PaaS-based cloud solution for e-Science

uDocker
PaaS Orchestrator
Infrastructure Manager
Identity and Access Manager

DEEP 1 platform
DEEP 2 platform

DEEP-Hybrid-DataCloud
AI/ML/DL PaaS, access to GPUs

AI4EOSC
Enhanced AI Platform for the #EOSC

eosc | FAIR-EASE
eosc | cancer
eosc | RAISE
eosc | Blue-Cloud2026
eosc | FAIR-IMPACT

Collaboration with INFRAEOSC projects

AI models will be integrated into AI4EOSC

FlexiGroBots
Fleets of autonomous robots for agriculture

iMagine platform, powered by AI4OS

iMagine (powered by AI4OS)
AI imaging platform for aquatic sciences

AI Compute platform for the EOSC Compute Platform

EGI-ACE
AI services for the EOSC Compute

Support to SME support in EOSC DIH

EOSC-Hub
Industry and innovate SME support

iMagine

For Operators

Platform features

Multi-site deployment

A service mesh approach

- Platform based on service mesh approach
  - Nomad: as workload management system
  - Consul: to enable service mesh
  - Traefik: to provide LB and proxy for running jobs
- AI4-PAPI (AI4 - Platform API)
  - Unified access to the platform, integrated with EGI AAI
  - Share-nothing architecture → horizontal scalability
- Automation of the deployment through Ansible roles
- Additional job (task) types to deliver more complex services
  - E.g. Image annotation, etc.
- AI4OS platform-wide container registry
  - Harbor https://registry.services.ai4os.eu/

iMagine

Traefik as service proxy

- Providing access to underlying tasks
- Rick monitoring of traffic status
- Multi-traefik deployment
  - 1 per Nomad datacenter
  - Pro: ensure access to user jobs in case of failure of other sites
  - Con: hostname will change if tasks are migrated
- Dynamic creation of endpoints (secure) for user tasks, e.g.:
  - IDE
  - API
  - Monitoring
  - Federated server

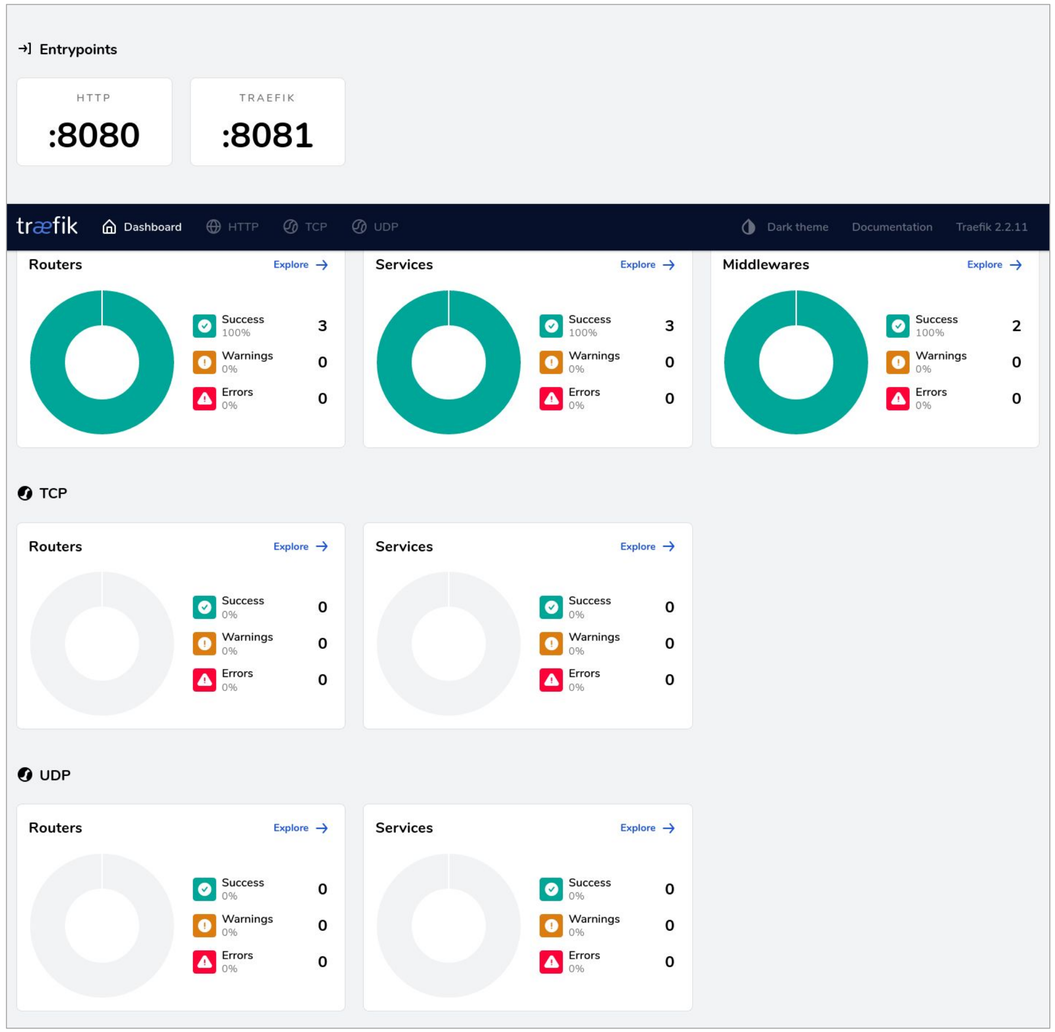

AI4-Platform API

Operator features

- Nomad provides very basic AuthN/Z system
- AI4-PAPI proxies user requests to the platform
  - i.e. no direct access to Nomad for the users
- Additional sidecar containers and tasks
  - Storage sidecar container
    - Currently supporting remote mounting through rclone
    - Integrating with EOSC-RAISE storage
  - Creation of metering tasks
    - Fine-grained accounting system
  - Execution of complex tooling deployments
    - E.g. image annotation, development environments, tracking services...
- Multi-API deployments are possible
  - Share-nothing architecture and stateless service

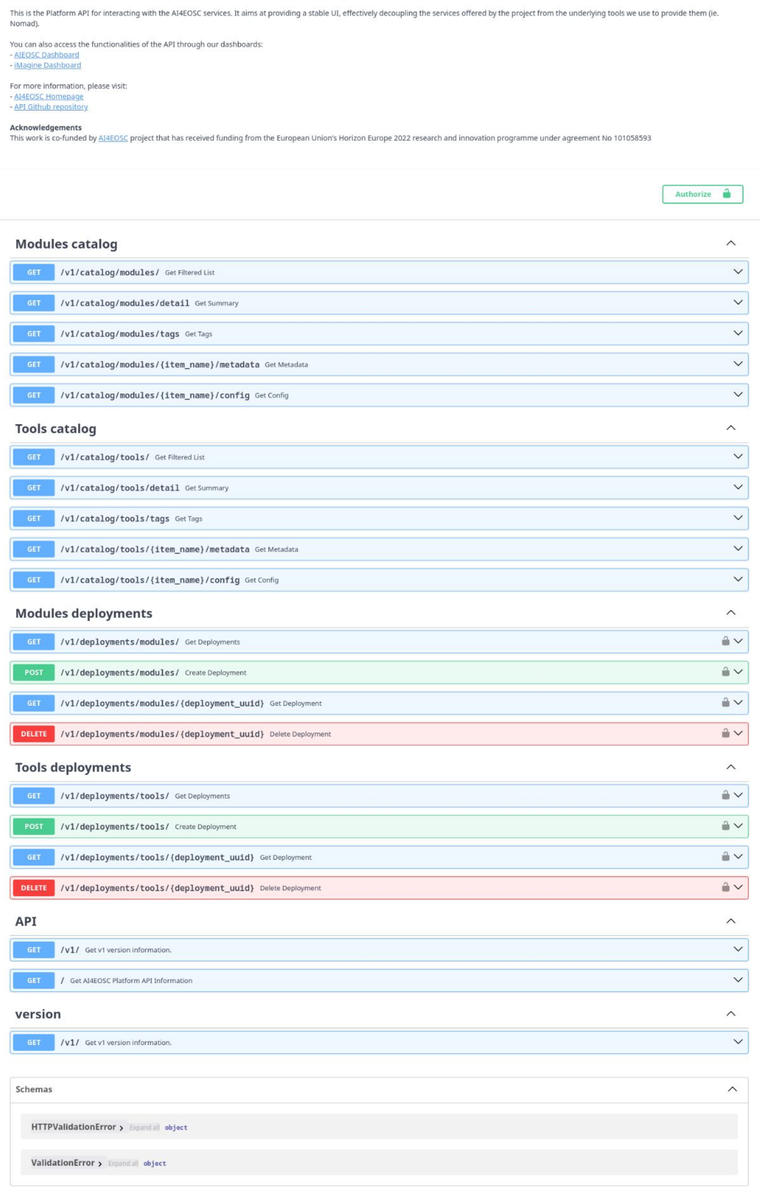

Platform next steps

Expected (operators) features for 2024

- Streamlining integration of new sites
  - Semi-automatically through Ansible (done)
  - Implement automation through IM (mid-term)

- Delivering iMagine AI Application Deployment Service
  - Initial work ongoing, but waiting for WP5 work to start
  - Different possibilities for deployment:
    - Through IM an a different cloud (done)
    - Through OSCAR in a platform-managed cluster (partially done)
    - Through OSCAR in user own resources (in progress)
    - Through AI4-PAPI in iMagine AI platform resources (planned)

iMagine

For Operators

Platform Architecture

AI4OS high level architecture

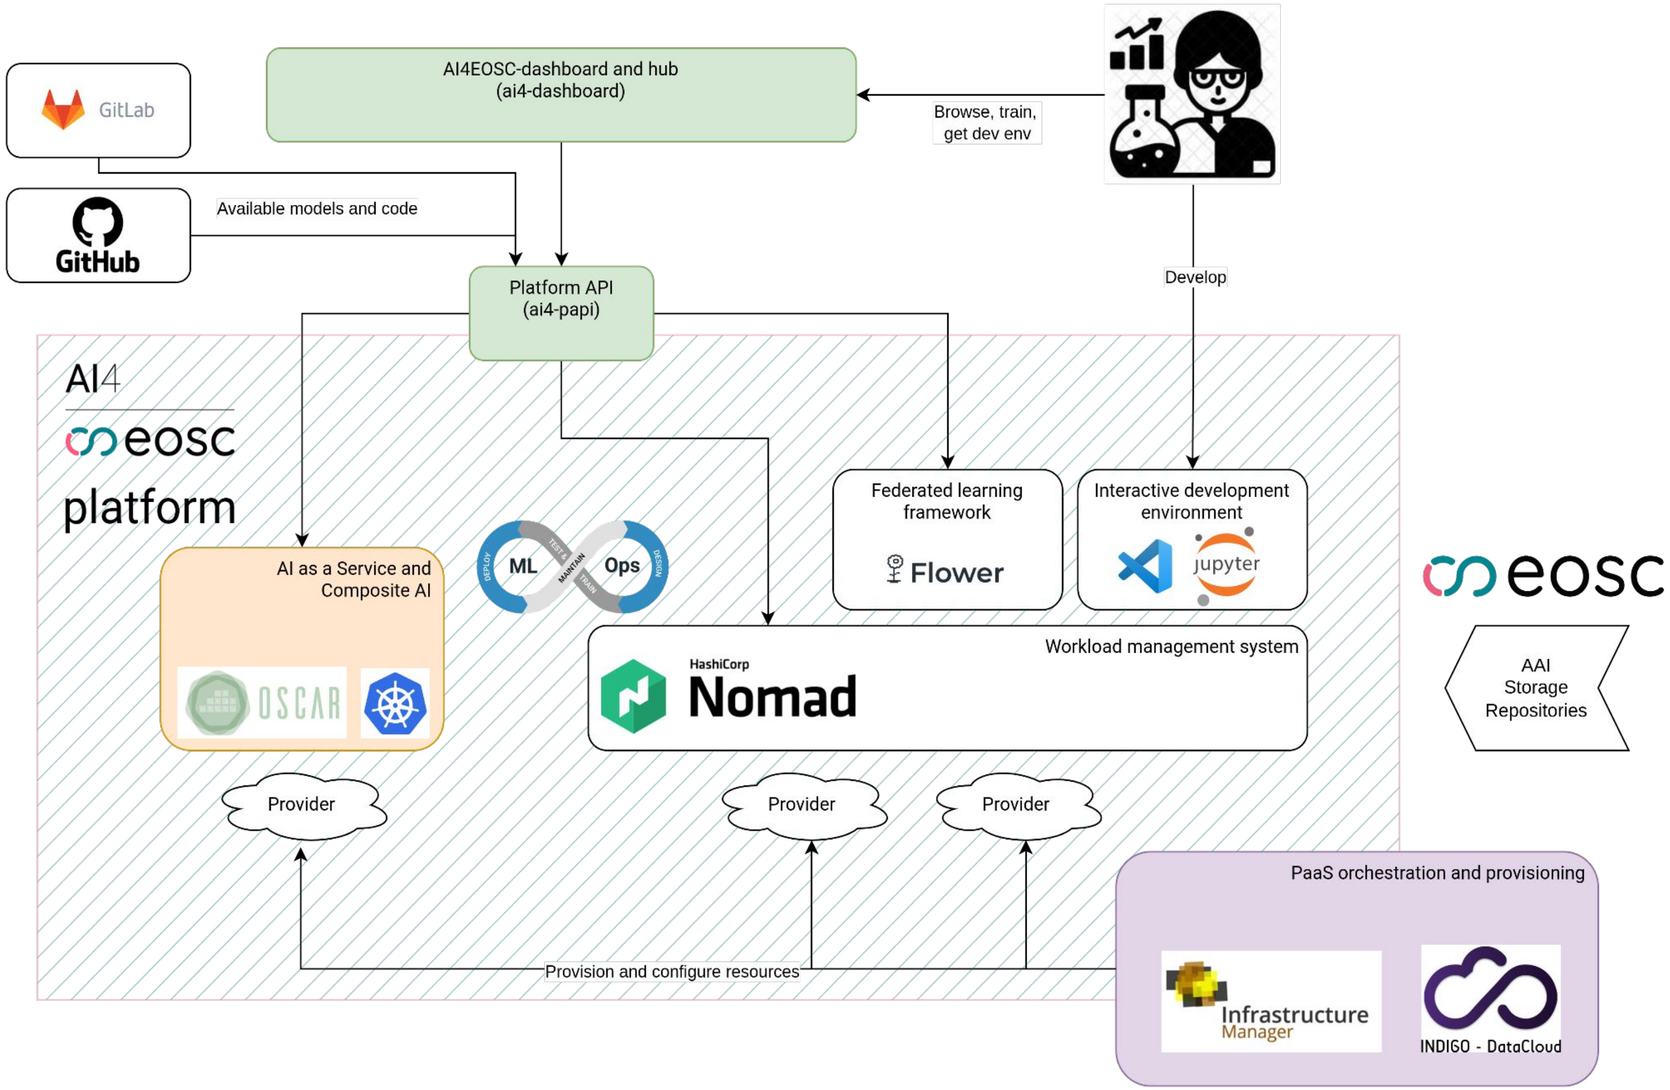

Detailed C4 architecture can be found here:

- Workspace https://structurizr.com/share/73873/2f769b91-f208-41b0-b79f-5e196435bdb1
- Diagrams: https://structurizr.com/share/73873/2f769b91-f208-41b0-b79f-5e196435bdb1/images

iMagine AI Platform system context (C4 Model)

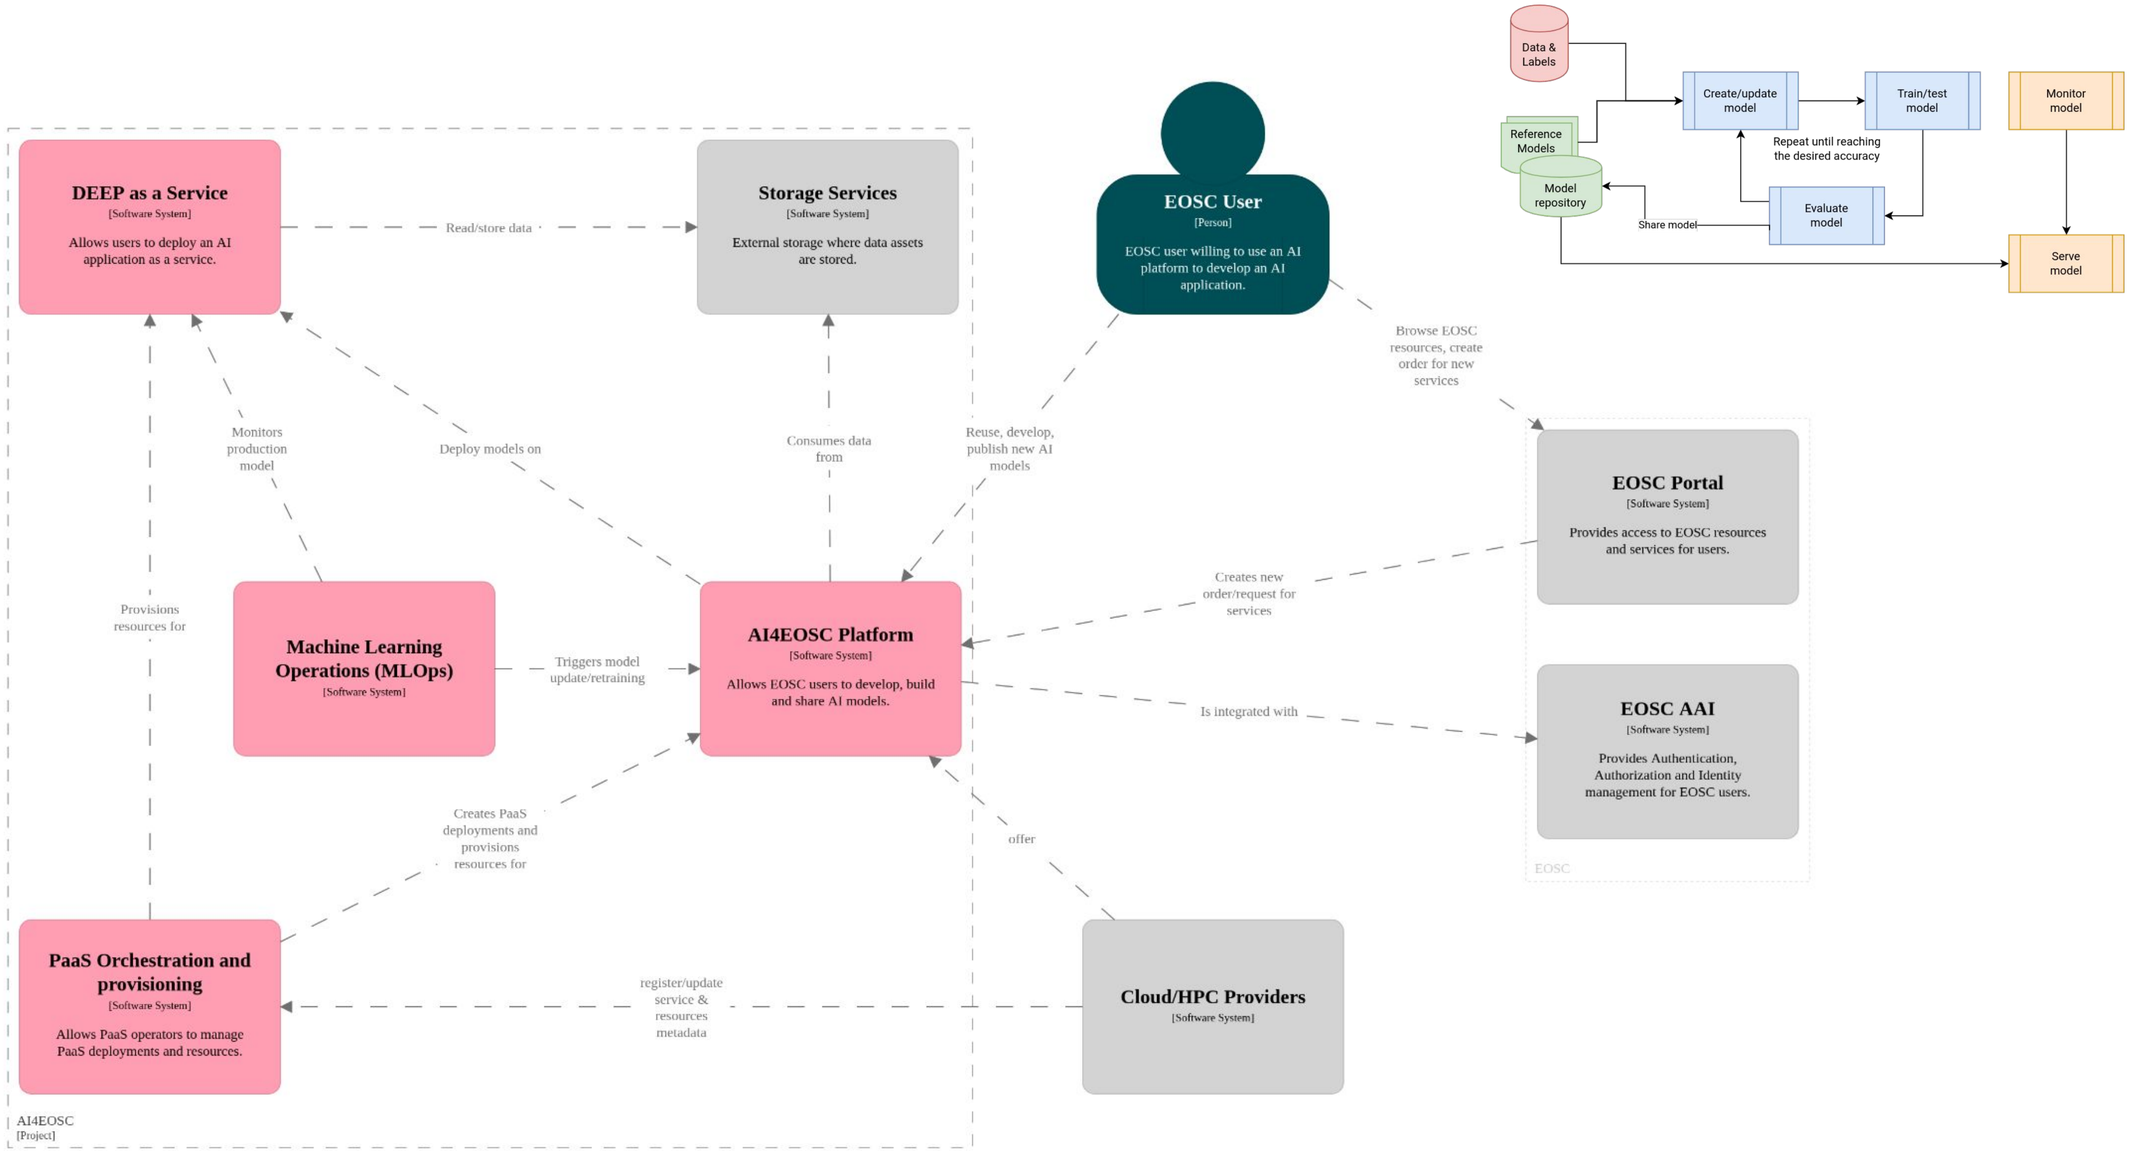

[System Context] AI4EOSC Platform

Tuesday, February 28, 2023 at 9:23 AM Coordinated Universal Time

iMagine

For users

Platform features

Services for AI/ML development

Training: upcoming features

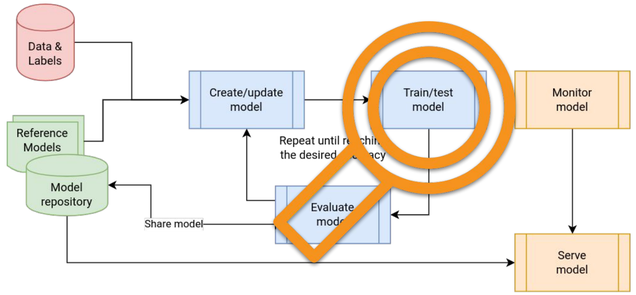

- Other distributed training schemes
  - Split learning
  - Parallel training with Horovod
  - TF Parameter server
- Model provenance system (MLFlow)
  - Including FAIR principles for ML models
- Experiment centric dashboard
  - i.e. group different models, data, trainings into one experiment
  - Allow to compare easily compare results
- Integration with additional online storage systems

Services for AI/ML development

Training: federated learning

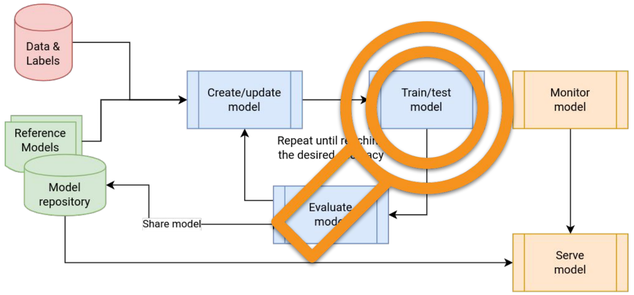

- Collaborative and decentralized approach to build ML models
  - No need to centralize a dataset (i.e. technical or privacy restrictions)
- Management of experiments through platform dashboard
- Participating clients both within AI4EOSC platform or external (with authentication)

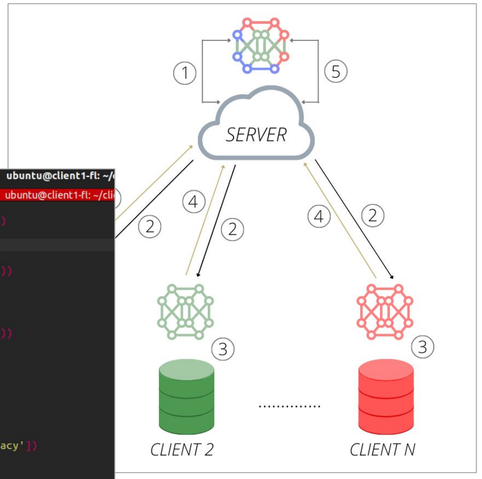

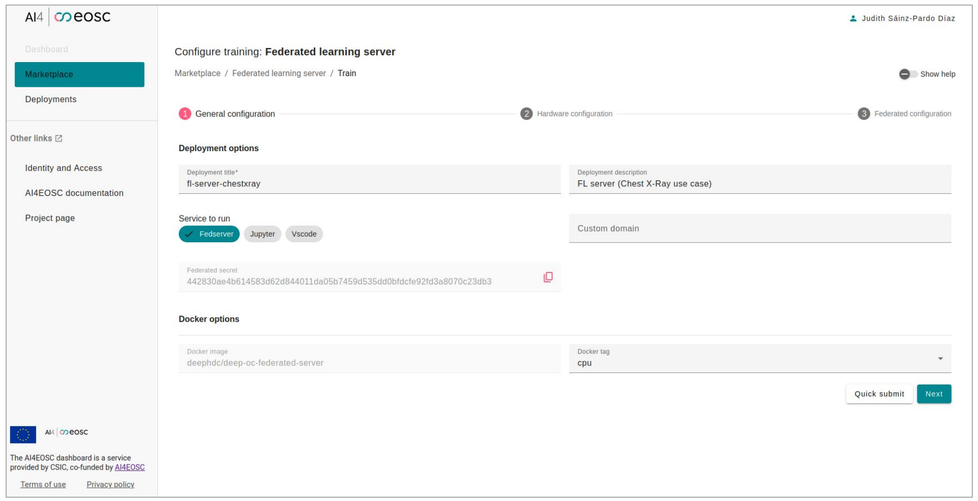

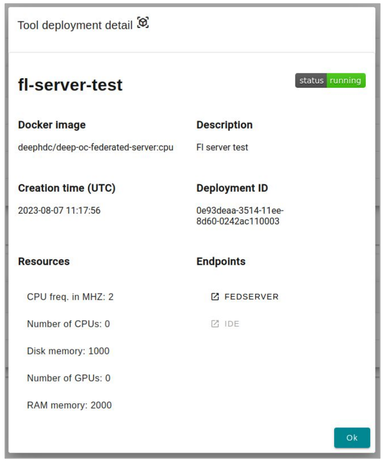

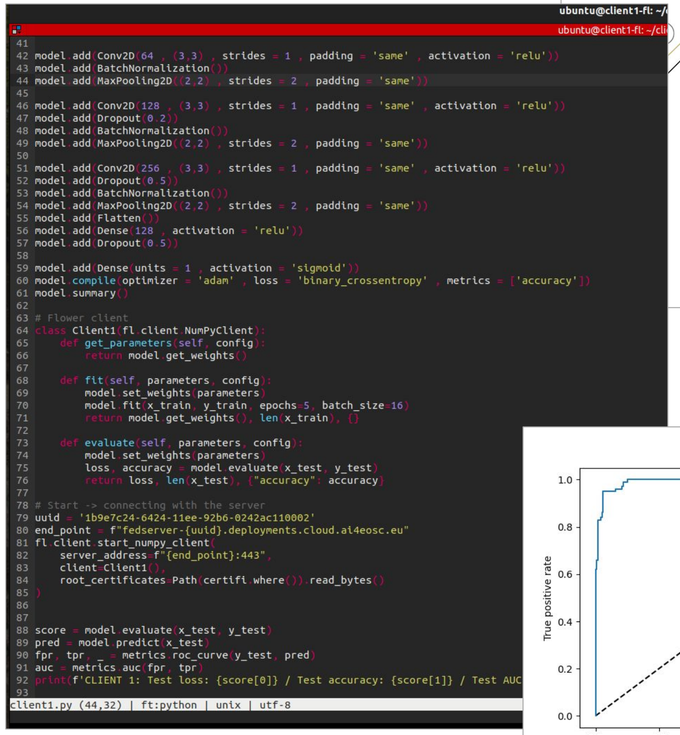

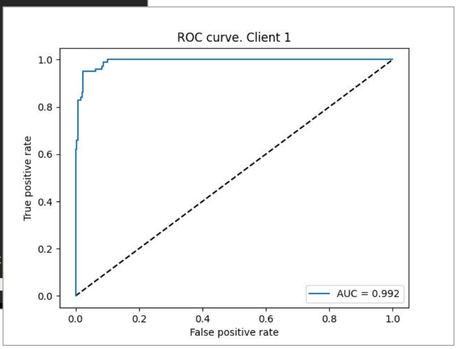

Services for AI/ML development

Deployment: OSCAR

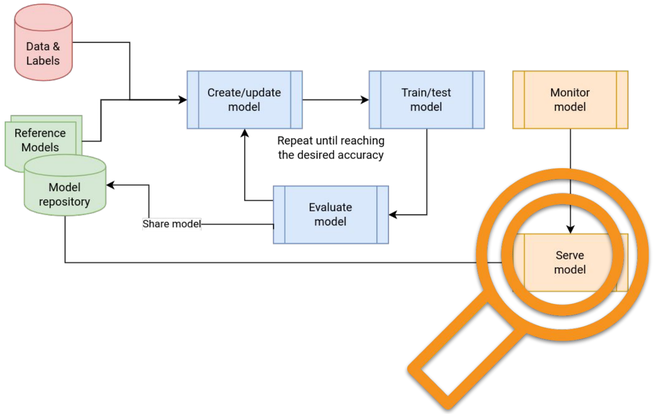

- OSCAR (https://oscar.grycap.net)
  to run the AI models for inference (AI as a Service)
  - Serverless event-driven execution
    - Asynchronous Mode: Files uploaded to the object-store trigger the invocation of a data-processing script that is run inside a container (out of user-defined Docker image) within a scalable Kubernetes cluster (e.g. batch jobs)

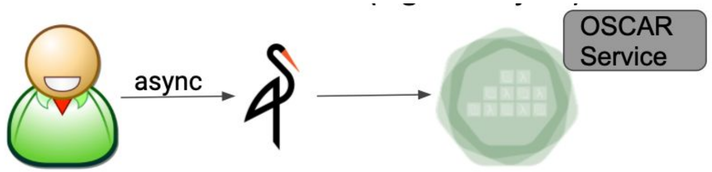

    - Synchronous mode: Scalable HTTP-based endpoints (based on KNative)

- https://inference.cloud.ai4eosc.eu/ui/
- https://inference.cloud.imagine-ai.eu/ui/

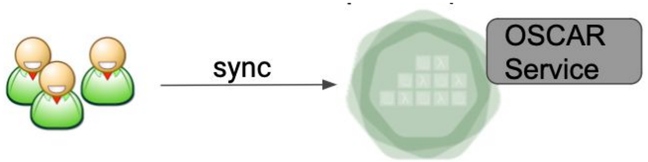

Services for AI/ML development

Deployment: Drift detection

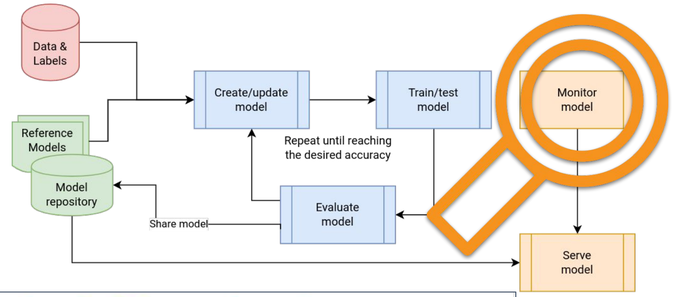

- Monitoring of models in production is not enough
  - Model learns from data, data is not stationary
  - Concept learnt by them model may change over time
- Data and concept drift detection → essential to build more robust models
- Frouros: state-of-the-art library for drift detection in ML problems
  - https://github.com/IFCA/frouros
- Ongoing work towards online services for drift detection → MLOps pipelines with drift detection

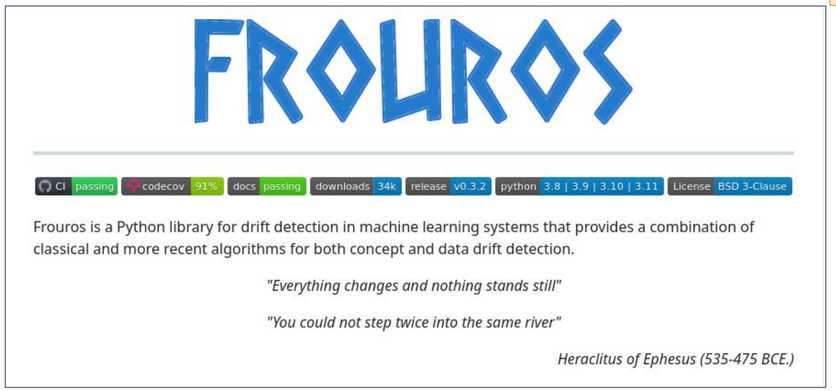

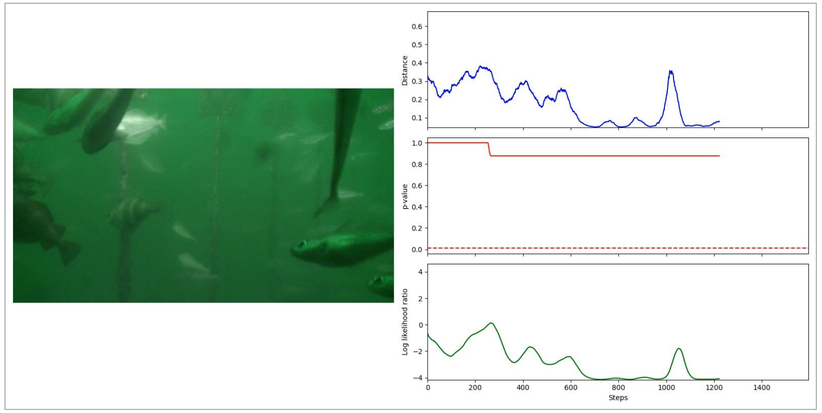

Example: data drift detection in underwater video

Services for AI/ML development

Deploy and monitor: upcoming

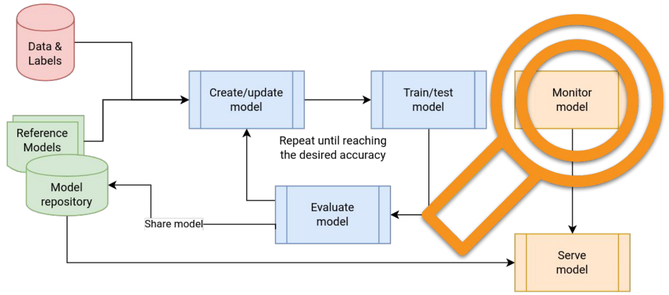

- Streamlining deployment and monitoring of models through CI/CD for ML applications and services
- MLOps pipelines definition
  - Automation for the whole ML lifecycle
  - Retraining based on given events (e.g. low accuracy)
- Compatibility with other ML deployment frameworks
- Improvements in API model definition (i.e. DEEPaaS)